*MULTIMEDIA EXTENSION # 1*

**International Journal of Robotics Research**

# Supplementary Information

for the paper

# *Low-bandwidth reflex-based control for lower power walking: 65 km on a single battery charge*

Most recent modification on January 17, 2014.


| | |
|---|---|
| Pranav A. Bhounsule | Disney Research, Pittsburgh, PA (formerly Cornell) |
| Jason Cortell | Mechanical Engineering, Cornell University |
| Anoop Grewal | Mechanical Engineering, Cornell University |
| Bram Hendriksen | Delft University of Technology, The Netherlands |
| J.G. Daniël Karssen, | Delft University of Technology, The Netherlands |
| Chandana Paul | Mechanical Engineering, Cornell; Applied Sciences, Harvard |
| Andy Ruina | Corresponding author: ruina@cornell.edu |
| | Mechanical Engineering, Cornell University, Ithaca NY |


This appendix provides some details about the robot Ranger, its hardware, software, simulation, optimization and the design of its controller. Please read the main paper before looking at this appendix. Even more details, including photos, videos and $100^+$ reports are on Cornell's www site:

http://ruina.tam.cornell.edu/research/topics/locomotion_and_robotics/ranger/Ranger2011/

## Contents

### The Machine


1. Introduction (and notation) ..... 2
2. Mechanical hardware ..... 6
3. Electronics ..... 14
4. Software on robot microcontrollers ..... 27


### Simulation and Control


5. Equations of motion ..... 41
6. Benchmark tests of the equations of motion ..... 49
7. System identification for mechanical parameters ..... 53
8. Modeling and system identification for motors and gearboxes ..... 59
9. Smoothings for simulations and optimizations ..... 66
10. Finite state machine for control of walking ..... 67

# 1 Introduction (and notation)

Ranger is a 4-legged bipedal robot built to be reliable while using little energy. Its four side-by-side legs are meant to mimic 2 legs. The outer pair of legs move together, acting as one leg, and so do the inner pair, hence the oxymoronic description '4-legged biped'. This robot lives, essentially, in two spatial dimensions (the sagittal plane). It has 3 main motors and 3 main internal degrees of freedom (hip and two ankles). Its culminating achievements are a 65 km walk and, later, the ability to walk with a total cost of transport of 0.19. Both of these seem to be bests for legged robots. This appendix describes some of the details of Ranger's construction and control.

| | |
|---|---|
| Total steps | 186,076 |
| Total time | 110,942 s (= 30 hrs 49 min 2 sec) |
| Number of laps | 307.75 |
| Lap distance | 212 m (= 0.132 miles) |
| Total distance | 65,243 m (= 65.24 km = 40.54 mi) |
| Average time per step | 0.6 s |
| Average distance per step | 0.35 m (= 13.78 in) |
| Average speed | 0.59 m/s (= 2.12 km/h = 1.32 mph) |
| Total power | 16 W |
| Power used by motors | 11.3 W |
| Power used by computers and sensors | 4.7 W |
| Total energy used | 493 watt-hours |
| Battery | 25.9 V Li-ion |
| Total robot mass | 9.91 kg (= 21.85 lb) |
| Battery mass | 2.8 kg (= 6.3 lb) |
| Total cost of transport (TCOT) | 0.28 (later lowered to 0.19) |

**Ranger's ultra-marathon walk.** On 1-2 May 2011 [9], Ranger walked non-stop for 40.5 miles (65 km) on Cornell's Barton Hall track without recharging or being touched by a human. Some of the crew that worked on Ranger are shown walking behind Ranger during the 65 km walk. Basic data are in the table above. Violeta Crow, at left, is steering.

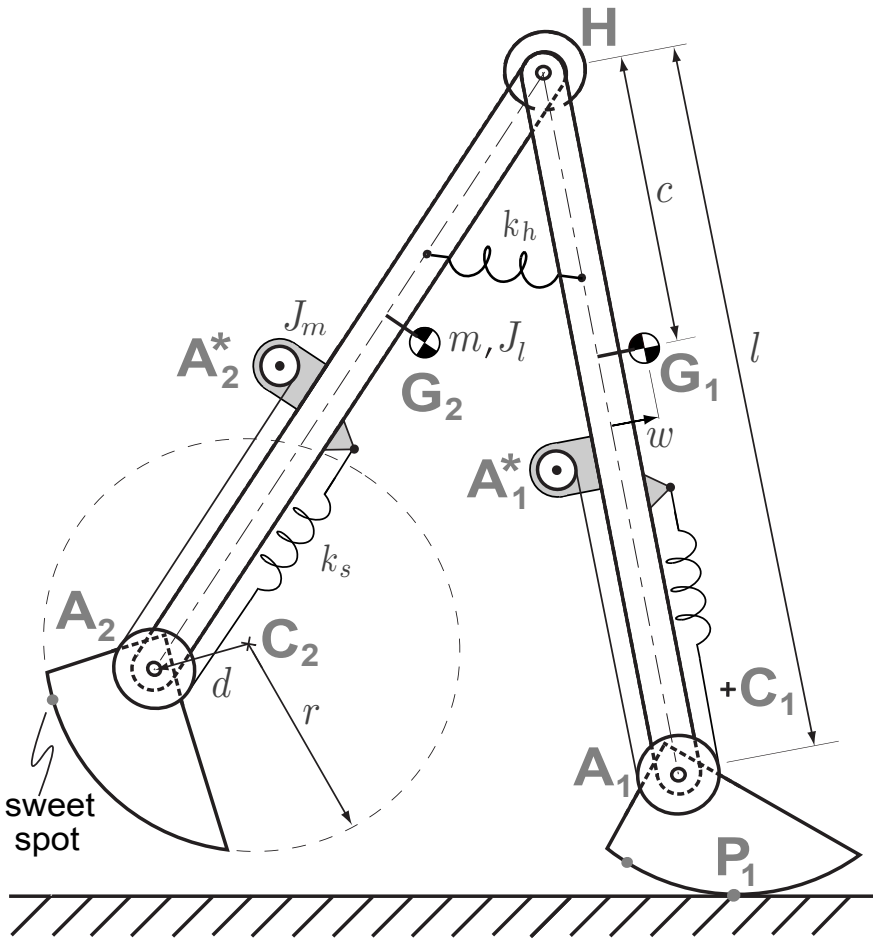


Figure 1: **Robot dimensions and feet geometry.** The feet bottoms are roughly circular arcs with radius $r$. The ankle joints $A_1$ and $A_2$ are offset from the center of circle by the distance $d$. As dictated by the geometry of circles the contact points $P_1$ and $P_2$ are always directly below the center of the circles $C_1$ and $C_2$, respectively, in level-ground walking. There is one foot configuration in which the ankle joint lies on the line joining the center of the circle and the contact point. For vertical ground forces this is a natural equilibrium position for the feet; it takes no ankle torque to hold the foot in this position. The contact point is then that part of the foot circular arc that is closest to the ankle. We call this point on the foot the 'sweet-spot'. The ankle motors are connected to the ankle joints via cables that we approximate as linear springs. The ankle motors ($A_1^*$, $A_2^*$) are actually nearly coincident with the hip H, but are separated in this diagram for clarity.

## 1.1 Notation

Variables used in this appendix are listed on the next two pages, some with respect to the side-view robot schematic figure 1. The main internal degrees of freedom are the hinges at the ankles ($A_1$, $A_2$) and hip (H). Because the ankle drive cables are elastic, the motors at $A_1^*$ and $A_2^*$ add two additional internal degrees of freedom.

### 1.1.1 Robot parameters

| Symbol | Value | Parameter description |
|---|---|---|
| $\ell$ | 0.96 m | Leg length. |
| $r$ | 0.2 m | Foot radius. |
| $d$ | 0.11 m | Ankle eccentricity. |
| $w$ | 0 | Fore-aft distance of COM. |
| $c$ | 0.15 m | Distance of COM from hip along the leg. |
| $k_h$ | 7.6 N m/rad | Hip Spring constant. |
| $k_s$ | 14 N m/rad | Ankle Spring constant. |
| $J_\ell$ | 0.45 kg m$^2$ | Inertia of legs about COM. |
| $J_{hip}$ | 0.55 kg m$^2$ | Inertia of legs about hip hinge. |
| $m$ | 4.96 kg | Mass of a leg. |
| $M_{tot}$ | 9.91 kg | Total robot mass. |
| $g$ | 9.81m/s$^2$ | Gravitational constant. |
| $\gamma$ | 0 | Ground slope (positive value is downhill). |
| $C_{1FW}$ | 0.05 N s/m | Coefficient of viscous friction between ground and stance leg. |
| $C_{2FW}$ | 0.05 N s/m | Coefficient of viscous friction between ground and trailing leg... in double stance. |

### 1.1.2 Motor parameters

| Symbol | Value | Parameter description |
|---|---|---|
| $G_H$ | 66 | Hip gear ratio. |
| $G_A$ | 34 | Ankle gear ratio. |
| $K$ | 0.018 N m/A | Motor torque constant. |
| $R$ | 1.3 Ω | Motor terminal resistance. |
| $V_c$ | 0.7 V | Contact voltage of the brush-commutator interface. |
| $J_m$ | 0.002 kg m$^2$ | Motor inertia. |
| $\mu_H$ | 0.1 | Coefficient of current dependent constant friction: hip motor. |
| $\mu_A$ | 0.1 | Coefficient of current dependent constant friction: ankle motor. |
| $C_{H1}$ | 0.01 N s/m | Coefficient of viscous friction: hip motor. |
| $C_{H0}$ | 0.1 N | Coefficient of constant friction: hip motor. |
| $C_{A1}$ | 0.01 N s/m | Coefficient of viscous friction: ankle motor. |
| $C_{A0}$ | 0.1 N | Coefficient of constant friction: ankle motor. |
| $P_{\text{fixed}}$ | 5.15 W | Power used by microprocessors, sensors and motor controller. |

### 1.1.3 Other variables

| Symbol | Variable description |
|---|---|
| $t$ | time. |
| $q_1$, $r_1$ | absolute angle made of stance foot wrt. vertical after and before heelstrike. |
| $q_2$, $r_2$ | relative angle between stance foot and stance leg after and before heelstrike. |
| $q_3$, $r_3$ | relative angle between legs; also hip angle after and before heelstrike. |
| $q_4$, $r_4$ | relative angle between swing foot and swing leg after and before heelstrike. |
| $q_{2m}$, $q_{4m}$ | motor angles at points $A_1^\star$ and $A_2^\star$ after heel-strike respectively. |
| $r_{2m}$, $r_{4m}$ | motor angles at points $A_1^\star$ and $A_2^\star$ before heel-strike respectively. |

| | |
|---|---|
| $x$, $y$ | world reference frame, $x$ in walking direction and $y$ is against gravity. |
| $x_h$, $y_h$ | $x$ and $y$ co-ordinate of the hip joint respectively. |
| $I_i$ | Motor current. $i = 2, 3, 4$ at the points $A_1^\star$, $H$ and $A_2^\star$ respectively. |
| $d_{\mathbf{step}}$ | Step length. |
| $t_{\mathbf{step}}$ | Step time. |
| $v_{\mathbf{step}}$ | Step velocity. |
| $M$ | Mass at the hip (for benchmarks in appendix 6). |
| $E$ | Energy. |
| $P$ | Power. |
| $N_i$ | Number of grid points, $i = ss$ (single stance) or $i = ds$ (double stance). |
| $\chi$ | Robot state vector and includes angles and angular rates. |
| $T_{1FW}$ | Torque between ground and stance foot. ($T_{1FW} = -C_{1FW}\dot{q}_1$). |
| $T_{2FW}$ | Torque between ground and trailing stance foot in double stance... ($T_{2FW} = -C_{2FW}(\dot{q}_1 + \dot{q}_2 - \dot{q}_3 - \dot{q}_4)$). |
| $T_3$ | Hip Motor Output Torque ($= G_H K I_3 - T_{fH}(I_3, \dot{q}_3)$) |
| $T_{fH}(I_3, \dot{q}_3)$ | Hip Motor Friction Torque ($= \mu_H \text{sgn}(\dot{q}_3) G_H K \lvert I_3 \rvert + C_{H1}\dot{q}_3 + C_{H0}\text{sgn}(\dot{q}_3)$). |
| $T_{fA}(I_i, \dot{q}_{im})$ | Ankle Motor Friction Torque ... ($= \mu_A \text{sgn}(\dot{q}_{im}) G_A K \lvert I_i \rvert + C_{A1}\dot{q}_{im} + C_{A0}\text{sgn}(\dot{q}_{im})$) where $i = 2, 4$. |
| $T_{2S}$ | Ankle Spring Torque ($T_{2S} = k_s(q_{2m} - q_2)$). |
| $T_{3S}$ | Hip Spring Torque ($T_{3S} = k_h q_3$). |
| $T_{4S}$ | Ankle Spring Torque ($T_{4S} = k_s(q_{4m} - q_4)$). |
| $F_{2S}$, $F_{2S}^{'}$ | Tensional force in the ankle cables at joint associated with dof. $q_2$. |
| $F_{4S}$, $F_{4S}^{'}$ | Tensional force in the ankle cables at joint associated with dof. $q_4$. |
| $H_i$, $V_i$ | Horizontal and vertical reaction forces respectively at joint $i$. |
| $H_i^\star$, $V_i^\star$ | Horizontal and vertical impulse respectively at joint $i$. |
| $\vec{g}$ | gravity vector ($= g\sin(\gamma)\hat{\imath} - g\cos(\gamma)\hat{\jmath}$). |
| $\vec{\omega}_1$ | Absolute angular velocity of $G_1$ after heelstrike ($\omega_1 = \dot{q}_1 + \dot{q}_2$). |
| $\vec{\omega}_2$ | Absolute angular velocity of $G_2$ after heelstrike ($\omega_2 = \dot{q}_1 + \dot{q}_2 - \dot{q}_3$). |
| $\vec{\omega'}_1$ | Absolute angular velocity of $G_1$ before heelstrike ($\omega_1' = \dot{r}_1 + \dot{r}_2$). |
| $\vec{\omega'}_2$ | Absolute angular velocity of $G_2$ before heelstrike ($\omega_2' = \dot{r}_1 + \dot{r}_2 - \dot{r}_3$). |
| $\vec{\alpha}_1$ | Absolute angular acceleration of $G_1$, the stance leg ($\alpha_1 = \ddot{q}_1 + \ddot{q}_2$). |
| $\vec{\alpha}_2$ | Absolute angular acceleration of $G_2$, the swing leg ($\alpha_2 = \ddot{q}_1 + \ddot{q}_2 - \ddot{q}_3$). |
| $\vec{v}_H$ | Velocity of point $H$ after heelstrike, ($= \dot{x}_h\hat{\imath} + \dot{y}_h\hat{\jmath}$). |
| $\vec{v}_{H'}$ | Velocity of point $H$ before heelstrike, ($= \dot{x}_{h'}\hat{\imath} + \dot{y}_{h'}\hat{\jmath}$). |
| $\vec{a}_H$ | Acceleration of point $H$ ($= \ddot{x}_h\hat{\imath} + \ddot{y}_h\hat{\jmath}$). |
| $\vec{v}_{G_1}$ | Velocity of point $G_1$ after heelstrike ($= \vec{v}_H + \vec{\omega}_1 \times \vec{r}_{G_1/H}$). |
| $\vec{v}_{G_1'}$ | Velocity of point $G_1$ before heelstrike ($= \vec{v}_{H'} + \vec{\omega}_1' \times \vec{r}_{G_1'/H'}$). |
| $\vec{a}_{G_1}$ | Acceleration of point $G_1$ ($= \vec{a}_H - \omega_1^2 \ \vec{r}_{G_1/H} + \vec{\alpha}_1 \times \vec{r}_{G_1/H}$). |
| $\vec{a}_{G_2}$ | Acceleration of point $G_2$ ($= \vec{a}_H - \omega_2^2 \ \vec{r}_{G_2/H} + \vec{\alpha}_2 \times \vec{r}_{G_2/H}$). |
| $\vec{v}_{G_2}$ | Velocity of point $G_2$ after heelstrike ($= \vec{v}_H + \vec{\omega}_2 \times \vec{r}_{G_2/H}$). |
| $\vec{v}_{G_2'}$ | Velocity of point $G_2$ before heelstrike ($= \vec{v}_{H'} + \vec{\omega}_2' \times \vec{r}_{G_2'/H'}$). |
| $\vec{r}_{H/P_1}$ | Position vector from point $P_1$ to point $H$ and so on. |

| | |
|---|---|
| $\vec{P}$ | Force on the trailing foot from ground in double stance ($= P_x\hat{\imath} + P_y\hat{\jmath}$). |
| $\vec{P^\star}$ | Impulse on the trailing foot from ground in double stance ($= P^\star_x\hat{\imath} + P^\star_y\hat{\jmath}$). |
| $\vec{H}_{/P_1}$ | Angular momentum about point $P_1$ and so on. |
| $\dot{\vec{H}}_{/P_1}$ | Rate of change of angular momentum about point $P_1$ and so on. |
| $\vec{M}_{/P_1}$ | External angular moment about point $P_1$ and so on. |
| $\|.\|$ | Absolute value function ($\|x\| = x$ for $x > 0$ and $\|x\| < -x$ for $x \leq 0$). |
| $[.]^+$ | Ramp function ($[x]^+ = x$ for $x > 0$ and $[x]^+ = 0$ for $x \leq 0$). |
| $\mathbf{sgn}(.)$ | Signum function ($\mathrm{sgn}(x) = x/\|x\|$). |

# 2 Mechanical hardware

Figure 2 shows the actual arrangement of Ranger's components. The two inner feet are linked directly by a hollow shaft, while the outer feet are driven by a single motor through separate cables that run through the truss above the hip hinge. Consistent with the facetious 'biped' moniker one can think of the outer leg (singular) and the inner leg. The steering is a fourth controlled degree of freedom: an electric motor twists the inner feet back and forth with each step, allowing Ranger to make gradual turns. The steering has little effect on the two-dimensional (sagittal plane) walking gait. The ankle joint motors have three main functions: pushing off at the end of stance, flipping up so the feet clear the ground during swing and keeping the feet close to flat on the ground during stance.

Most of the robot mass is near the hip. With light legs and light feet, leg swing can be quick using small torque. Further, because of the light legs, the stance leg dynamics become nearly decoupled from the swing leg dynamics. Ranger's legs are hollow carbon fiber tubes and the feet are cut from high-strength 7075 aluminum alloy stock. The ankle joint motors are located on the hip axis.

Structurally, Ranger has three 'thighs': one aluminum box supporting the inner pair of legs, and one more for each of the outer two legs. The outer leg boxes are spanned by an aluminum top truss, which also supports electrical wires and ankle-drive cables. Thus the three thighs are effectively two, an outer and an inner thigh. Protection against the occasional fall is provided by protruding foam 'eye' stalks and foam 'ears'.

The robot has a total mass of 9.91 kg and a leg length of 1 m, measured from the bottom of the foot up to the hip axis. The total height is about 1.1 meters, not including antennae.

## 2.1 Hip drive

Hip actuation is via a 46-watt-nominal brushed direct current motor (Faulhaber #2657W012CR) and a 66:1 planetary gearhead. The hip motor and gearhead are mounted in the outer leg box, aligned axially with the hip joint. The gearhead output shaft is attached to the inner leg via a flexible coupling (flexble in bending, stiff for shaft rotation), to prevent bearing loads due to slight construction misalignment between the hip bearing mounts and the motor mount.

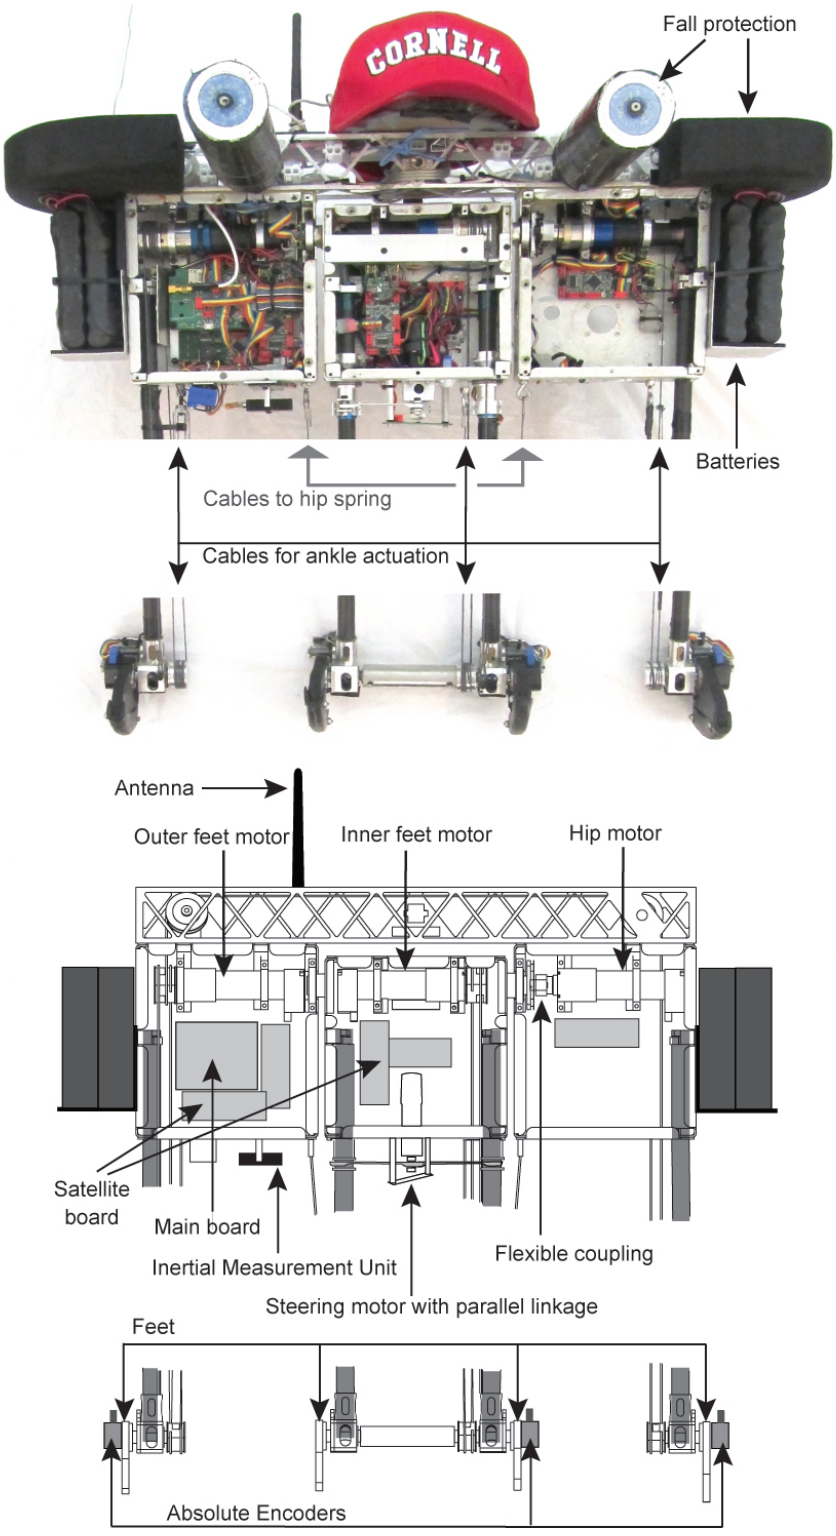


Figure 2: **Overall layout, front view.** In the photograph the boxes are open. The foam assemblies for fall protection ('eyes' and 'ears') serve no other purpose. The hat is only decorative. The main and satellite boards contain most of the electronics. More details are in the text and following illustrations.

## 2.2 Ankle drive system

Each pair of feet is actuated by a 46-watt-nominal brushed direct current motor (Faulhaber #2657W012CR) with a 43:1 planetary gearhead.

**Cables and springs.** The motors drive the feet through a system of thin stainless steel cables (figure 3). The cable at the back of the ankle joint (analogous to the Achilles tendon in humans) goes directly from the foot pulley to the motor pulley. However, there is no direct connection between the front of the foot pulley and the front of the motor pulley. Instead, the two cable ends coming off the fronts of the pulleys are attached to a pair of springs, which are in turn attached to fixed points on the leg assembly. The spring at the foot pulley causes the foot to flip upward when the ankle motor relaxes the (Achilles tendon) drive cable. The spring at the motor pulley balances the force of the ankle spring, so the motor need not waste power continually fighting it. The motor spring also prevents slack when the motor is off. The pre-stressed springs, chosen for a small spring constant and a large initial extension, behave almost like constant force springs, allowing the motor to move the ankle through its full range of motion with minimal interference from the springs.

The inner feet are rigidly connected so that they can be actuated by one cable. For the outer feet two cables are used, wound in tandem on a single motor pulley (see figure 3).

The ankle-drive cables need to be small diameter to reduce cable fatigue when used over small-diameter pulleys. We used 7 x 19 strand core low-stretch nylon-coated stainless steel cable made by Sava Industries. The cable for the inner feet (model 2050SN2) had an overall (with coating) diameter of 1.59 mm, an uncoated (stainless-only) diameter of 1.19 mm, and a breaking strength of 1200 N. The outer feet used model 2037SN cable, with a coated diameter of 1.17 mm, an uncoated diameter of 0.96 mm, and a breaking strength of 710 N. These cables had non-negligible compliance, stretching under tension and thus, with the motors, making up series-elastic actuators. Two lengths of cable are used to drive the outer ankle pair, but only one for the inner ankle pair (see figure 3). To ensure symmetry of control and behavior for both pairs of ankles, we needed to appropriately match all three cable compliances as follows:

**Matching cable compliances.** For a laterally symmetric gait the cables between the outer ankle motor and the ankles are in parallel (e.g. for a small rotation of the outer ankle motor, with the ankles joints held fixed, each cable stretches an equal distance). Thus the outer cable's spring constants add. For gait symmetry the single cable driving the inner ankles should have the same spring constant as the combined value from the outer ankle cables, so the inner cable should be twice as stiff as each of the outer-ankle cables. To double the stiffness, a larger cable with double the cross-sectional area was used for the inner feet. Further, the cable to the left outer ankle needs to go up and across the top truss before heading down to the ankle, making the left cable about twice as long as the right outer ankle cable, and giving it twice the compliance. To make the left and right Achilles tendons match, the left cable was threaded through a length of thin carbon fiber tubing running almost the full length of the leg and glued securely, thus reducing the stretchable portion of the cable length.

**Attachments.** Stress concentrations at the connections of the tendon cables with the carbon fiber tubes were reduced by tapering the tube thickness at those points, and by leaving the nylon cable coating in place for several centimeters into the tube. Farther inside the tube, the nylon was stripped off the cable, allowing a strong epoxy bond to the stainless steel strands within. In an

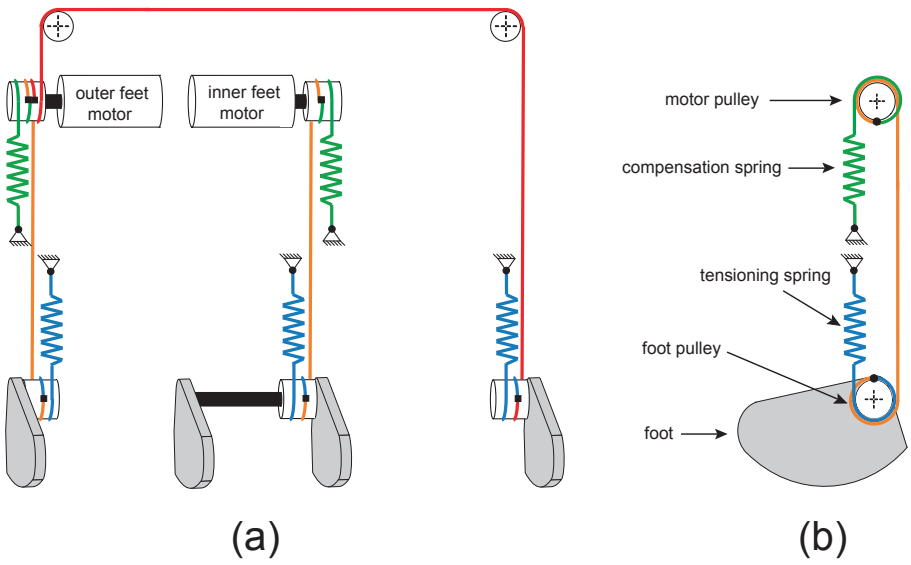


Figure 3: **Ankle actuation. (a)** Perspective schematic, **(b)** Side schematic of one leg. One motor controls the inner ankles and one controls the outer ankles. For each leg pair ankle extension (e.g. push-off) is powered by a single cable and retraction (for foot ground clearance) is powered by a return spring. Another return spring on the motor prevents cable slack when the motor is off, and also balances out the force of the ankle return spring so that the motor doesn't have to.

early incarnation steel was used instead of carbon-fiber and the connection with the cable was with solder. One robot failure was from tendon cable breakage due to corrosion from acidic flux residues from this soldering.

The cable ends are clamped inside the pulleys.

Because there are angle sensors at both ends of the drive cables (at the ankle joints and also at the drive motor shaft), the ankle torque can be measured by multiplying the cable stretch by the cable stiffness. In the end, however, we made no use of this torque measurement in the controller.

### 2.3 Ankle joints

There is some subtlety in the ankle design (figure 4(a)). The ankle needs to be light and strong while also providing:

1. a secure connection to the carbon fiber leg tube,
2. a low-friction rotational joint with minimal play,
3. a stop to limit the motion of the foot,
4. and a path to take the foot contact sensor cable out of the foot and up the leg tube.

while holding up to hundreds of thousands of heel-strike impacts.

The 7075 aluminum ankle shaft is partially hollow and has an opening in the middle so that the sensor cable can go from the foot up into the hollow leg tube. The sensor cable is wrapped

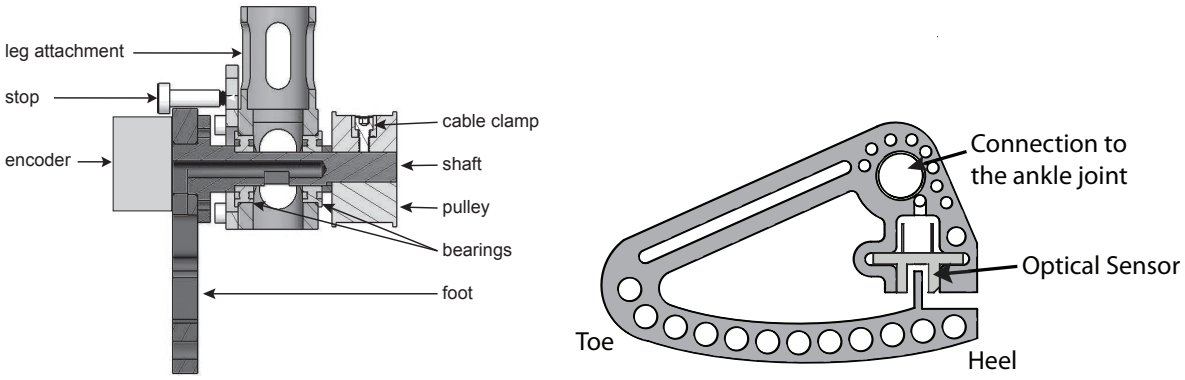


(a) Cross-section of the ankle

(b) Foot side view

Figure 4: **Ankle and foot. a)** Front view of a cross-section of the foot-ankle assembly, **b)** Side view of the foot. The top of the foot hits the stop when the top of the foot is approximately vertical. The 'stop' includes a spring-loaded electrical contact, thus acting as both a mechanical stop and an electrical limit switch.

around the foot shaft once or twice, spiraling outward. Thus, when the shaft rotates, the resulting bend in the cable is distributed over a substantial length, minimizing local strains and maximizing flexure life. The cable used was selected for long life in high-flex environments, with extra-fine-gauge conductors made of high-strength copper alloys. (e.g., 3M high-flex ribbon cable, series 3319.)

Each ankle shaft is supported by two ball bearings, flanged for ease of fabrication. The foot screws to a flange machined in one piece with the shaft. The opposite end of the ankle shaft has a flat machined into it, forming a "D" in cross section. The foot pulley clamps on with a matching D-shaped opening, thus preventing slippage and play. Initial student designs used just a set screw or just a clamp to connect the pulley to the foot-drive shaft. Such arrangements were found to be unreliable with too much slip or play or both.

## 2.4 Feet

The foot shape is shown in figure 4(b). Its bottom edge is a portion of a larger circular arc centered above and in front of the ankle joint (point $C_2$ in figure 1 on page 3). The foot is shaped and positioned such that when made to stand in its equilibrium position on flat ground the point of foot contact (the sweet spot) is part of the circular arc. This point is near the heel. It takes zero torque to maintain the foot in this position when the ground reaction is vertical.

The feet are also load cells cut in an open C shape. A vertical load on the bottom of the foot closes the gap at the heel. The size of this gap is measured with an optical sensor. Finite element analysis was used to design the foot so that expected loads would cause deformation within the range of the optical sensor. Further, the gap at the back was sized to close for large loads, thus preventing plastic deformation of the foot from unusually large heel-strike forces.

## 2.5 Support structure (hips/thighs)

The hips/thigh for the outer legs are constructed of two folded and spot-welded 5052 aluminum sheet boxes, with a top truss (machined from 6061-T6 aluminum square tubing) connecting the two

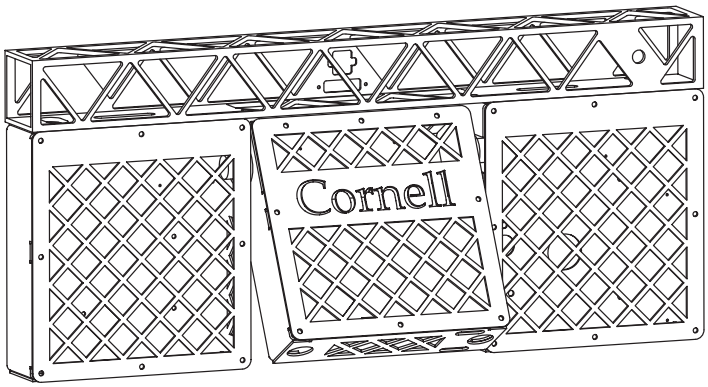


Figure 5: **Legs supporting structures (hips/thighs).** Each outer leg has one closed aluminum box made from folded and welded sheet stock. The pair of inner legs has a similar closed box. The outer boxes are joined by a truss made from from a hollow aluminum box beam with triangles machined out to reduce weight. The inner box is hinged to the outer boxes by the hip axle which runs parallel to, and below, the truss.

boxes. The inner leg is constructed of one box (figure 5). Fully-closed box structures were chosen for their light weight, high stiffness, high strength, and ease of fabrication. The motor mounts and top truss were bonded with Mereco Metregrip 303 epoxy adhesive, with screws as a backup. The rectangular boxes were also convenient for enclosing electronics, motors, batteries, etc. Because the box stiffness and strength are reduced by orders of magnitude when the cover plates are removed, the robot cannot operate without the box covers.

For stationary bench testing of the hip and outer ankle motors, the top bar provides sufficient support to resist motor torques even with the covers off, but the inner box does not have such support. For testing with the inner box cover removed an angle bracket was fastened across the motor mounts, to provide some supplemental support.

### 2.6 Hip spring

The robot has a pair of tension springs connecting the inner legs with the outer legs. One end of each spring is connected to the inner legs, near the bottom, while the opposite end is connected to the outer boxes. The spring length is minimum when the legs are parallel; therefore, because they are pre-stretched, they effectively act as a torsion spring at the hip joint, pulling the legs towards a parallel configuration (see figure 6b). Using figure 6, we can calculate the torque at the hip as

$$
\begin{aligned}
T_{\text{hip}} &= 2kab\left(1 - \frac{\ell_0}{\ell}\right)\sin(\theta) \\
&= 2kab\left(1 - \frac{\ell_0}{\sqrt{a^2 + b^2 - 2ab\cos(\theta)}}\right)\sin(\theta) \qquad (1) \\
&\approx \underbrace{2kab\left(1 - \frac{\ell_0}{b-a}\right)}_{k_h}\theta \equiv k_h\theta \qquad (2)
\end{aligned}
$$

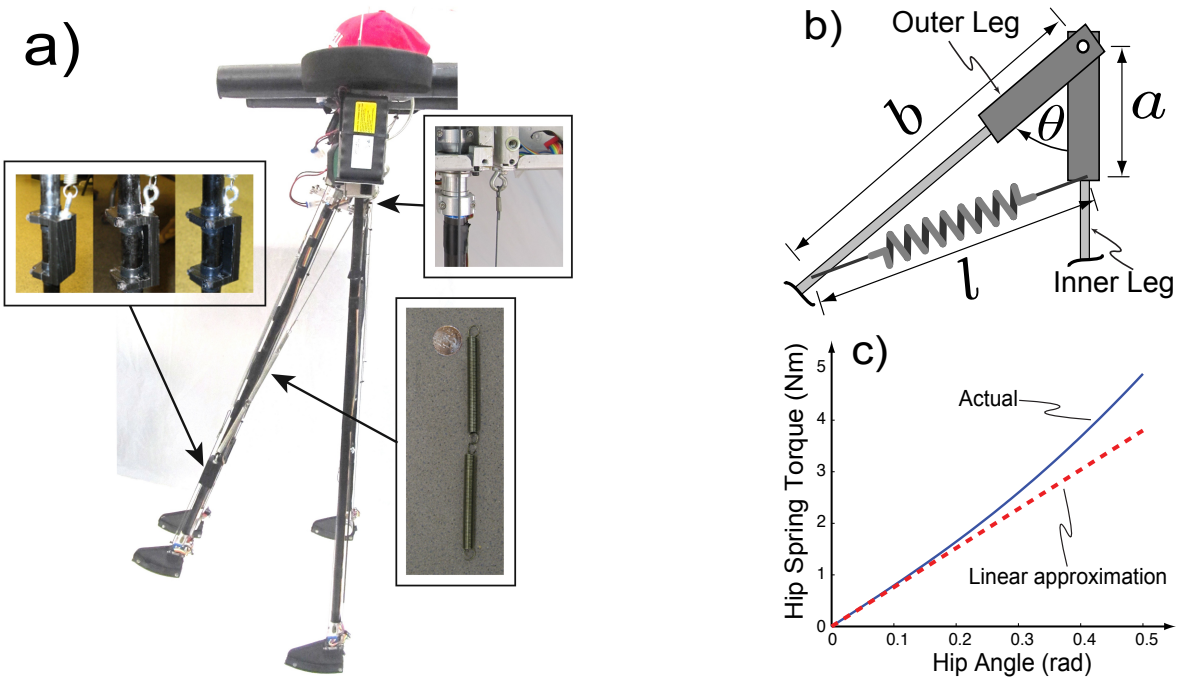


Figure 6: **Hip Spring.** The hip spring acts as a torsional spring tending to keep the outer legs parallel with the inner legs. Its purpose is to speed the leg swing with less motor effort. **a)** a side-view photograph showing also the spring and attachments. **b)** Schematic diagram of the hip spring, **c)** Plot of hip spring torque vs hip angle. Blue solid line is hip spring torque calculated using equation 1. The linear approximation is equation 2, which we use in simulations. The nearly linear behavior follows from the spring pre-stretch, which gives nearly constant tension.

where $T_{\text{hip}}$ is hip spring torque, $k$ is the spring constant of the tension spring, $a$ and $b$ are distances from the hinge joint to the connection point on the inner and outer leg respectively, $\ell_0$ and $\ell$ are the lengths of the spring when the legs are parallel and when the hip is at an angle $\theta$, respectively. $k_h$ is the small-angle effective spring constant. The spring stiffness $k$ was chosen so that the hip motor barely had the strength (6 N m) to hold its leg at an angle of 0.4 rad, fighting both the spring and the weight of the swing leg.

### 2.7 Steering

Ranger can steer around gradual curves via a twisting degree of freedom in the inner legs (see figure 7). A small 1-watt-nominal Maxon gearmotor is connected to the inner legs via a pair of linkages, and controlled by one of the ARM7 satellite boards. For steering, the two inner leg tubes are rotated along their axis inside bearings in the inner leg box. The two inner ankles are rigidly attached to each other at the bottom, and the leg tubes themselves are quite thin and relatively flexible in bending. Thus, rotation of the two leg tubes around their own axes at the top, via the motor-driven linkage, results in rotation of both inner feet around a common central vertical axis at the bottom. By thus swiveling the feet, the robot can turn by up to about 5 degrees in one step of the inner legs. Coordination of inner leg twist with leg swing happens in the top-level processor; however, the direction and amount of this steering has thus far been controlled by a

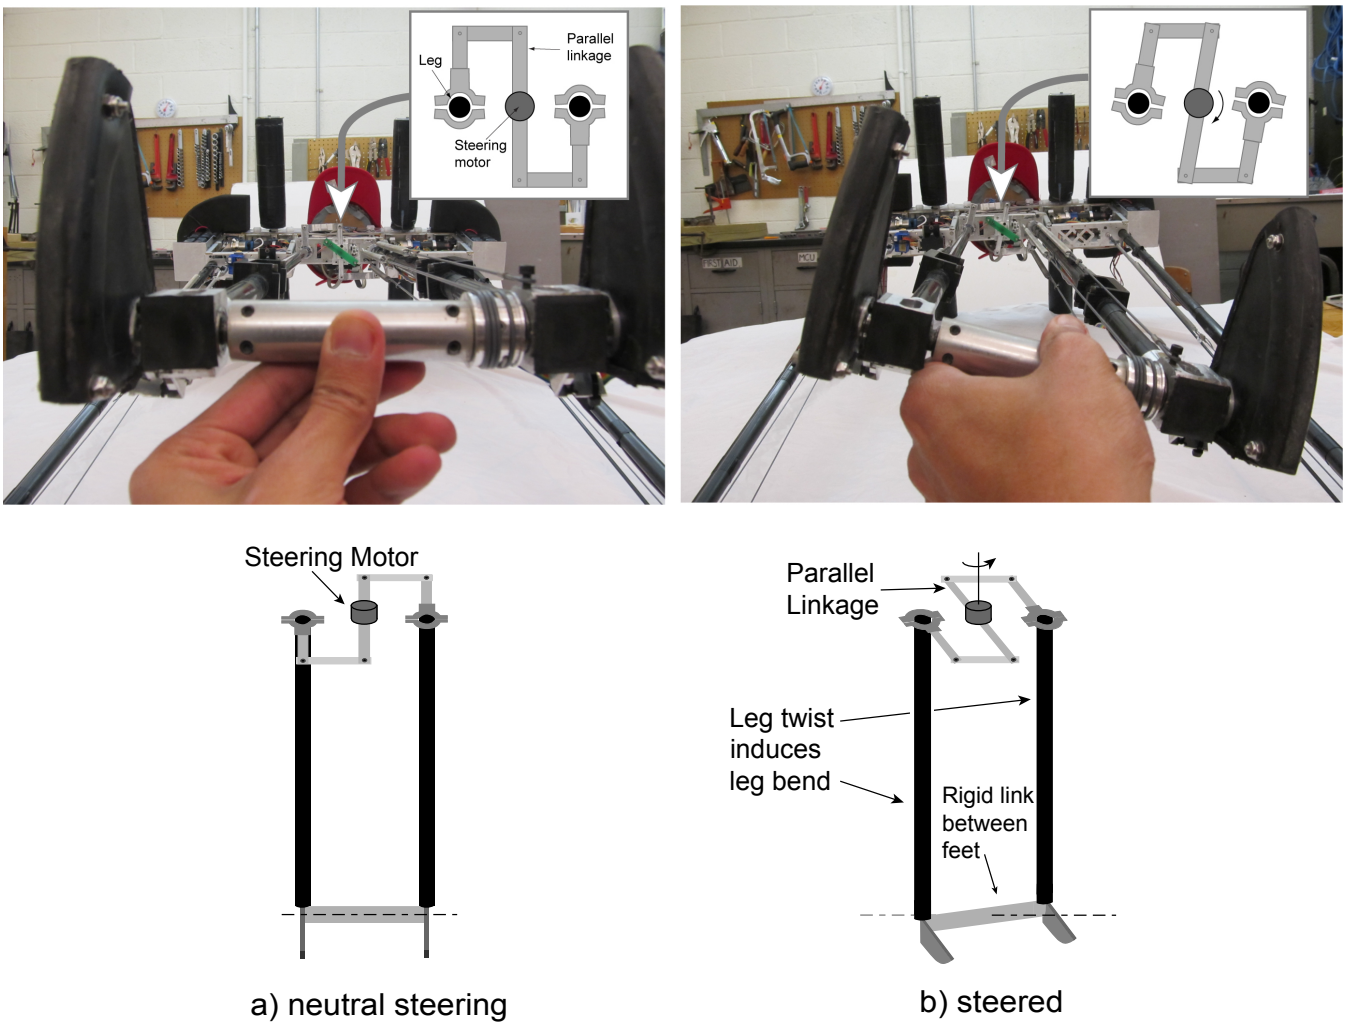


Figure 7: **Steering mechanism.** Steering is by twist of the inner legs. To turn the robot to the right the inner legs are turned left during the inner-leg stance phase; the feet do not slide on the ground so the robot turns. Neutral steering is shown on the left, twisted and bent legs are shown on the right. The steering motor turns the parallel linkage which twists the legs (the legs are mounted at the hips in bearings that allow axial turning). Because of the rigid link between the ankles this twisting bends the legs resulting in a turning of the pair of feet.

human-operated remote control.

**Previous steering mechanism.** The mechanism described above replaces a previous ill-conceived mechanism. We had reasoned (and observed in an earlier robot) that any break in the robot's lateral symmetry would result in a turn one way or the other. So we had a motorized lever with a pulley that adjusted the length of the left outer ankle drive cable, thus changing the angle of the left foot relative to the right. This mechanism did result in some steering, but the result was highly sensitive in amount and even direction to gait, slope, and flooring material. In retrospect, the design was something like trying to steer a car by deflating the tires on one side; here we were trying to steer

a robot by making it stub one or the other of its toes on the floor with each step. This steering was used for the 1 km and 9 km walks, but, for example, the 1 km walk ended abruptly when the robot was unintentionally steered over a sand pit, and the 9 km walk nearly ended with the robot heading out the back door of the gym and into the street.

### 2.8 Fall protection

The fore-aft cylinders with painted-on eyes and the foam ears (see figure 2) are for shock absorption in case of falls. The eye stalks protect against falls directly forward or backwards. Polyethylene foam pipe insulation/pool floats are guided by nylon tubes passing through holes in the top truss of the robot, and held in place by plastic caps on the ends. During a fall, the nylon tubes slide through and allow the foam to compress, absorbing the energy of the fall, but keep the foam from buckling or bending. The design was tested with a virtual robot made of wood. The final design cushioned forwards or backwards falls over a stop distance of many centimeters. Through the life of Ranger these 'eyes' have cushioned on the order of 20 falls. In early trials of any controller the user holds a string that prevents falls all the way to the ground.

The "ears" are high-density foam rubber, intended to protect the robot in the unlikely event of a sideways or twisting fall. (As a 2D biped, Ranger so far has always fallen directly forward or backward).

## 3 Electronics hardware

### 3.1 Evolution

Ranger's electronic hardware, its so-called nervous system, went through several iterations. The first electronics version (named the KoBrain, after its student designer Ko Ihara) used a custom board with a Freescale MC56F8347 hybrid DSP microcontroller, external A/D converter, and three-axis MEMS accelerometers and rate gyros. A plug-in daughter board held three motor driver ICs, a wireless data communications module, and various sensor interface circuitry. The KoBrain was used for the first long distance walk attempt (1 km). At that point we realized the value of having angle sensors at all the joints, in addition to the motor shaft angle, and also wanted to install a commercial IMU (Inertial Measurement Unit). The Freescale microcontroller, though selected for its large input-output capability, was nevertheless running low on ports and processor time by this point. So three additional MC56F805 microcontroller boards were added, one per joint, communicating with the main one via CAN bus. These were bulky and not particularly energy-efficient. The third robot nervous system, described in detail below, was a network of ARM-based processor boards designed for a future biped robot with ten or more degrees of freedom. This system was installed and debugged on Ranger.

### 3.2 Robot nervous system overview

The CAN (controller area network) bus was introduced in the automotive industry in the late 1980s to accommodate the increasing numbers of sensors and actuators in each vehicle. It allowed manufacturers to install multiple electronic control units (ECU) throughout vehicles, each handling a local subset of the sensor and control load, and communicating with each other over the

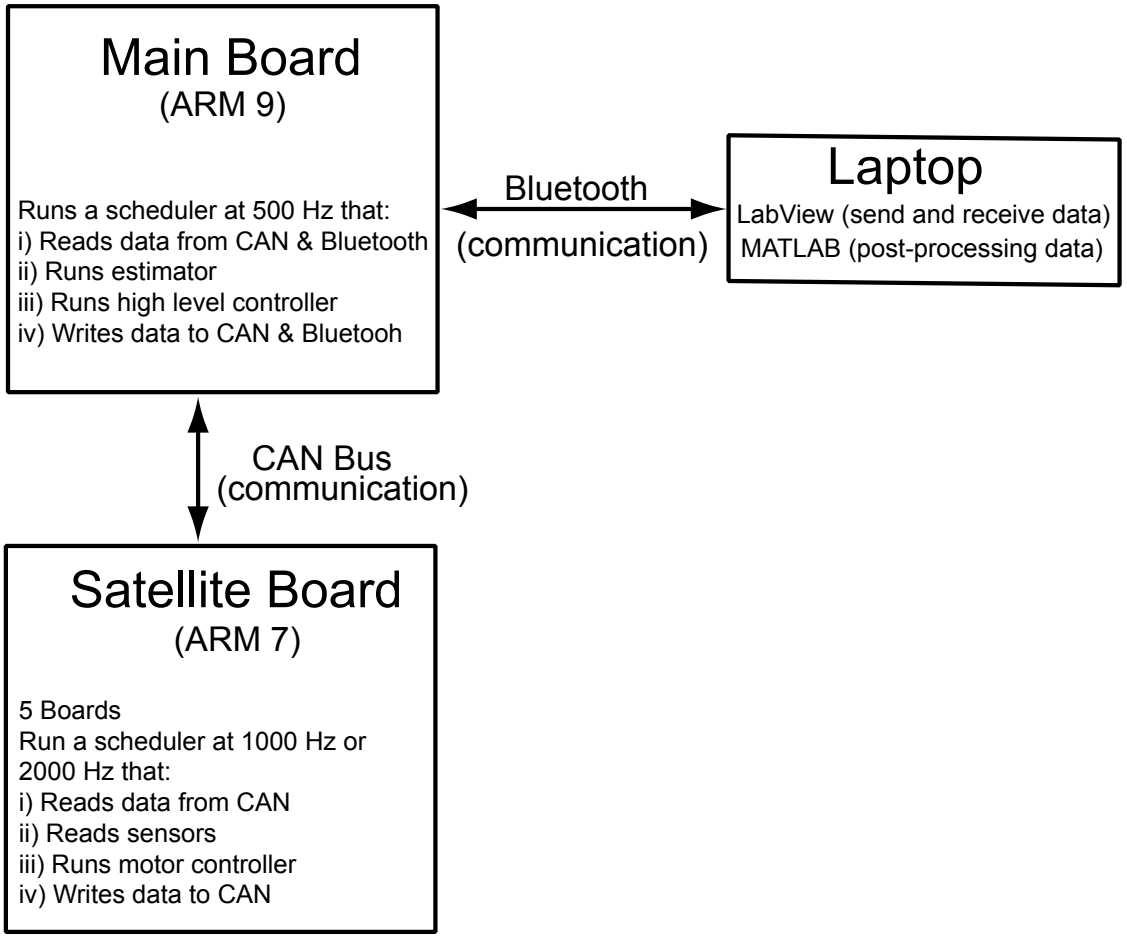


Figure 8: **Overall electronics architecture.** An ARM 9 'main brain' communicates with 6 ARM 7 satellite processors via a CAN network. Data collection and tuning use a Bluetooth wireless data link. Not shown is the hobby radio control used for steering.

high-reliability CAN network. This greatly reduced the size and cost of the wiring harnesses required; for example, masses of costly shielded low-level sensor cables could be replaced with a single twisted-pair CAN cable and a power cable. CAN also increased the modularity of the automotive electronics: ECUs from different manufacturers could be added or removed from the system with minimal impact on the rest of the electronics and support software. The needs for a robot are similar. Ranger uses a version of the automotive CAN bus, but with a cable and network protocol modified to work better with our small, low-power robot.

Ranger's ECUs are divided into two classes: top-level "main brain" processors based on the ARM9, and sensor-and-actuator satellite ("spinal") processors based on the ARM7. The satellite boards are connected to four high-speed CAN buses; connections between buses and to one or more top-level processors are made via a satellite board set up as a CAN router (see figure 8). The boards were designed for low energy use, and the modular design allows boards, sensors, and motors to be added or removed quite easily. Wiring complexity is reduced, since the boards are

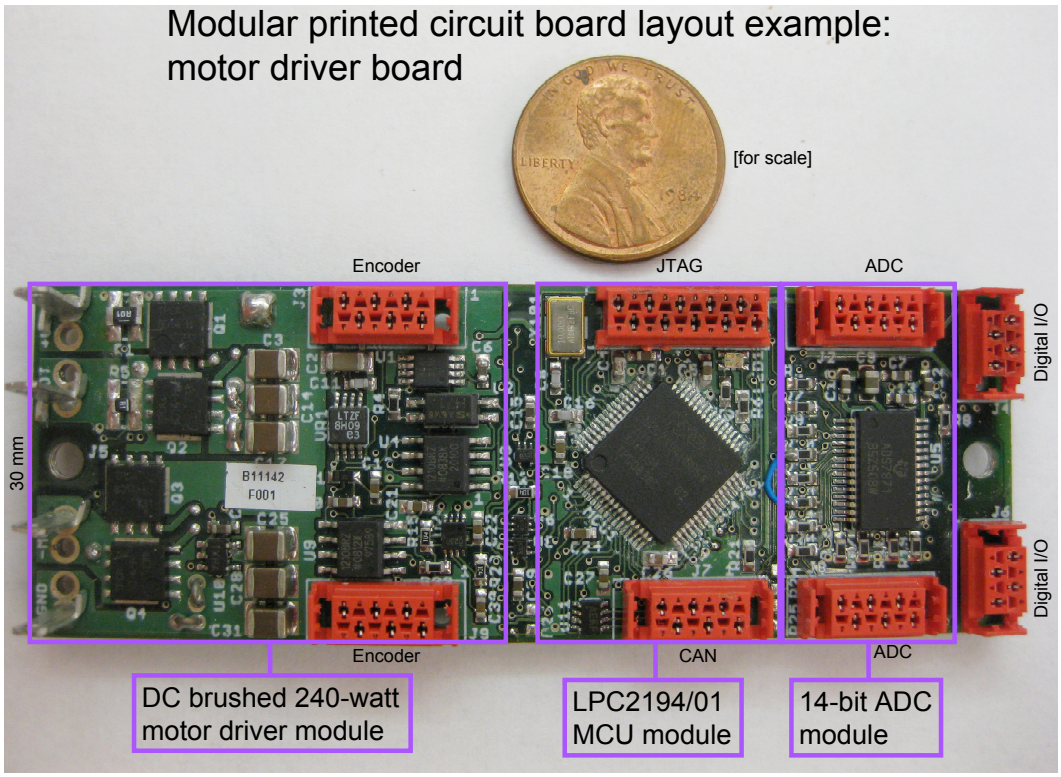


Figure 9: **Example satellite processor board.**

connected to each other only by power supply and CAN bus cables.

### 3.3 Top-level (main brain) processor

Ranger uses a Phytec phyCORE LPC3250 system-on-module for running high-level control and estimation. With an NXP 208 MHz ARM 9 processor and 32 MB of memory on the module, it is clearly less computationally capable than a modern PC, but also uses only 370 mW. The LPC3250 was selected because at the time it was one of the only low-power processors available with floating point capability, which we have deemed essential for our top-level controller. Earlier designs with only fixed-point calculations used up substantial programming time with managing arithmetic. At the time of design, available PowerPC modules used more than three watts for comparable performance, and PC-104 boards with floating point used over ten watts. The LPC3250 performance has been adequate, although has demanded some extra attention to computational efficiency.

The network architecture allows additional top-level processors to be added as needed; for example, a second processor could be (but hasn't yet been) added to handle sensor fusion and state estimation. In addition, other low-power boards with floating point have become available now and could be used to upgrade the top-level processor, such as the tiny Overo modules from Gumstix Computing (based on TI OMAP processors), and a variety of Intel Atom boards consuming about 5 W but giving netbook performance levels.

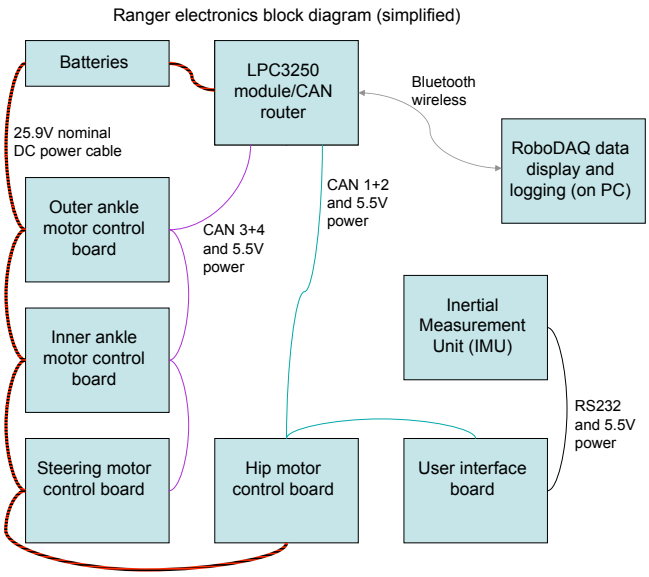
Ranger electronics block diagram (simplified)
Batteries
LPC3250 module/CAN router
Bluetooth wireless
RoboDAQ data display and logging (on PC)
25.9V nominal DC power cable
Outer ankle motor control board
CAN 3+4 and 5.5V power
CAN 1+2 and 5.5V power
Inertial Measurement Unit (IMU)
Inner ankle motor control board
RS232 and 5.5V power
Steering motor control board
Hip motor control board
User interface board

Figure 10: **Wiring between microprocessor boards.** Dark red is main power. Violet and light blue are CAN buses. The CAN router board is the port for information to and from the main brain (which is not shown).

### 3.4 Satellite board processor

For the ARM 7 satellite boards we selected the NXP LPC2194/01 microcontroller (MCU), mostly on the basis of its low power consumption, about 80 mW at its maximum 60 MHz clock speed. Keeping the power per processor low was critical if we were to keep the power requirements of a whole network of them within reasonable bounds (Note that in the end, even with our attention to low power electronics, half of Ranger's energy budget was used for processing and sensing). A second requirement was support for multiple CAN buses. The LPC2194 is identical to the better-known LPC2129, but with an additional two CAN controllers. Other peripheral support was adequate, though certainly not as extensive as the previous Freescale processors. As with the top-level processor, new higher-performance MCUs are now available that would meet the satellite board requirements, including the ARM Cortex M3 LPC17xx series from NXP and the Cortex M4-based Freescale Kinetis series.

### 3.5 Modular satellite board design

As per the network design philosophy, the satellite boards needed to be customized where necessary to the particular sensors and actuators they were expected to handle. We wanted a modular satellite board design, to allow reuse of electronics and code between the satellites whenever possible. After some initial work on a stacking board system, as used earlier on Ranger, we realized that we used far too much board area and design effort in implementing the stacking connectors. It appeared instead that we would do much better to design modules at the printed circuit board (PCB) layout stage of board development. Functional modules, for example the circuitry for the MCU and an external analog-to-digital converter (ADC), could be designed and optimized into small board

layout blocks, then combined and configured to give the desired satellite board capabilities. No space is thus wasted on interboard connectors, and the flexibility of the process allows for far more diversity in satellite board design. Our PCB layout CAD program, EAGLE, does allow layout art and its associated schematics to be copied and pasted between boards, albeit awkwardly, so we used it to develop a set of layout modules. These are standardized at 30 mm height, and are intended to be placed side-by-side to form a satellite board (although they can also be dropped as needed onto larger boards).

### 3.6 PCB layout modules used in Ranger

An example module is shown in figure 9. The various modules used are listed here.

**LPC2194/01 MCU module:** Includes the MCU and other circuitry common to all the satellite boards: a pair of CAN transceivers, three voltage regulators (5 V, 3.3 V, and 1.8 V, with the last two switching for higher efficiency), JTAG programming connector, precision 10 MHz crystal, power-on reset circuit, tri-color LED, and CAN network/power connector, all on a 20 x 30 mm 4-layer board. In addition, one board layer is largely devoted to a set of parallel traces that make available all the peripheral pins of the MCU at both sides of the board, allowing connections as needed to other modules on either side.

**ADC module:** Based on the TI ADS7871 8-channel 14-bit analog-to-digital converter (ADC), the ADC module has a voltage reference and RC input filters, and is 15 mm long.

**DC motor controller module:** The brushed DC motor driver has a MOSFET H-bridge with optical isolation; sensors for battery voltage, battery current, motor current, and board temperature; two encoder inputs; a watchdog timer (to shut off the motors if the motor control software crashes); and can drive a motor continuously at 30 V, 8 A, and 100 kHz. (Maximum pulse-width-modulation (PWM) frequency is 200 kHz.) The motor control module is 40 mm long.

**System power module:** Power to each board is delivered via 5.5 V on an extra pair of wires in the CAN network cable, and stepped down from the battery voltage by a switching regulator. The system power module also has sensors for measuring current and voltage input from the battery. Only the CAN router board uses the system power module and this powers the sensors and processors on all the other boards. Motor power comes directly from the batteries, using separate, larger-gauge wires.

**Optional modules:** The layout module concept is flexible enough to allow other, larger board sizes and miscellaneous other circuitry to be added where needed, while still being able to reuse the core module layouts.

### 3.7 Satellite boards used in Ranger

The wiring layout for Ranger, as a whole, is shown in figure 10. Here is a list of the different boards used, based on the modules above.

**B2A Brushed DC Motor Control:** Each of Ranger's four motors is connected to one of these boards, which uses a motor controller module, MCU module, and ADC module placed in a

line with a little space in between to allow for interconnection. The overall size is 85 x 30 mm, and it uses under 200 mW, not including the motor and MOSFETs; the MOSFET drive uses another 300 mW at 30 V and 100 kHz when active.

**B8A User Interface:** The User Interface (UI) board includes the MCU module with a 2-line LCD display, multicolor LCD driver, radio control inputs, pushbutton switches, an RS232 port, six color LEDs, and a piezoelectric speaker. The board size is about 100 x 100 mm.

**B10A CAN Router/Main Brain Carrier/Bluetooth:** The CAN router board includes the MCU module, the system power supply, four CAN ports, a Roving Networks RN-21 Bluetooth wireless module, a MicroSD port, a USB port, an EtherNet port, and connectors for plugging in the LPC3250 system module. It is the same size as the LPC3250 module, 58 x 70 mm.

### 3.8 Ranger sensors and interface hardware

Ranger has about 40 sensors and various other peripheral devices either on, or monitored by, the various satellite boards. Ideally each sensor should be on or very near a satellite board; the boards would then condition the potentially noise-prone raw sensor signal and convert it to floating-point readings transmitted onto the CAN bus. Keeping the boards in close proximity to their peripherals would reduce electrical noise problems and also reduce wiring. But because the modular electronic system was a retrofit for Ranger, we decided not to put boards at the feet, and instead ran the sensor cables up the legs. A couple of additional sensor cables cross over the top bar to connect to the logically closest board, rather than the physically closest. For example, the hip angle sensor is in the right outer leg box, but its cable runs to the hip motor control board in the left outer leg box, next to the hip motor. This was intended to simplify the programming. Unfortunately this separation also resulted, not surprisingly, in some problems with electrical noise.

Similarly, the motor control boards should be located as closely as possible to the motors. This reduces the electrical interference they cause, reduces resistive losses, and reduces the overall bulk and complexity of the wiring. In Ranger the motors are indeed quite close to their controllers. Ideally, however, the controller would be part of the motor assembly.

**Motor encoders:** Each of the three main motors is equipped with an incremental encoder on the motor shaft. The direct connection allows for simpler control at high gains. The hip motor uses a HEDS 5500 optical encoder, with 192 pulses per revolution and consuming about 100 mW. The two ankle motors use IE2 - 512 magnetic encoders, with 512 counts per revolution; they use about 50 mW each. The driver software uses timer capture inputs on the LPC2194/01 to detect transitions at the quadrature outputs from the encoders, time the transitions, and increment or decrement the position count as required by the motor rotational direction.

**Angle encoders:** Ranger has magnetic absolute angle sensors on the inner ankle assembly, both outer ankle joints, the hip joint, and the steering linkage. Except for the steering, these are based on the RLS AM8192B Hall array angle sensors, as used in their RMK3 evaluation modules. A small magnet is glued into a hole in the center of a non-magnetic shaft, and the RMK3 module is mounted over the magnet, with a small gap. The Hall sensor array detects rotation of the magnet's field, and these sensor readings are converted into a digital representation of the joint angle by the circuits in the AM8192B. The satellite board reads in encoder information in two forms:

1. Via the SSI (serial output) port on the AM8192B, which is read in by an SPI (serial peripheral interface) port on the LPC2194/01 as a binary number representing the absolute angle of the joint.
2. The AM8192B also has standard quadrature-format incremental encoder outputs, which are read in by the timer capture inputs as described above. The quadrature outputs allow an estimated joint angular velocity to be found, along with a largely redundant relative angle reading.

On Ranger, the signals for both of these interfaces pass through a single unshielded 6-wire ribbon cable to the appropriate satellite board. RLS recommended using differential line drivers for the AM8192B outputs, but as a demo board the RMK3 didn't come with any and we didn't install any. It turned out that the output of the SSI interface in particular had quite high impedance and was thus noise-prone, despite being a digital signal. The AM8192B sensors worked well otherwise. But at about 200 mW each the 4 angle encoders make up a significant fraction of the power budget for the electronics (0.8 of 5 watts). The left outer ankle sensor and its associated satellite board were deemed redundant and, to save power, disabled for the 65 km walk.

The steering linkage uses an Asahi EM-3242 magnetic angle sensor, which has a 10-bit-resolution analog output connected to one of the ADC inputs on a satellite board.

**Battery current sensors:** We used the Maxim MAX4081S for its small size, low power, high gain (60 V/V), and high-voltage (76 V) capability. The MAX4081S IC allows bidirectional measurement of current by amplifying the voltage across a low-value current sense resistor, but can only be used on the high (positive) side of a power circuit. On Ranger the MAX4081S was used to measure the battery current going into or out of the motor control board, with good results. The similar unidirectional MAX4080S IC was used to measure current flow into the CAN router board, which powers all of the electronics except for the motors and motor drivers. The current sensor outputs go to a differential input pair on the ADS7871 ADC (or to one of the LPC2194/01's internal ADC inputs, in the case of the CAN router board.) We used current sense resistors with solid metal elements (Vishay/Dale WSLP series) because we found that the film on metal film resistors vaporized when subjected to capacitor inrush currents.

**Motor current sensors:** The motor current sensor attempts to find the current going to the motor leads by measuring the voltage drop across a low-value current sense resistor installed in series with the motor, amplifying this voltage, and then sending it to the ADS7871 ADC, which converts it to a digital value proportional to current. When the motor lead with the sense resistor is on the fixed-voltage (non-PWM) side of the motor driver H-bridge, measuring the sense voltage is not an unusually challenging task. However, when the other half of the H-bridge is driven by the PWM signal, the current sense resistor is subject to a 100 kHz AC signal at the full battery voltage. The part selected for this job, the Analog Devices AD8210, does derive a fairly clean current signal from such a noisy environment. Unfortunately, the AD8210's zero-current reading is not the same for the PWM signal as it is for DC, giving a discontinuity in the motor current readings as the PWM values pass through zero. We tried an Allegro ACS712 Hall effect sensor instead, and it worked great on the bench, but when we installed it in Ranger we saw that it wouldn't work properly when sitting next to the

powerful permanent magnets of a DC motor. So instead we wrote an odd bit of code which switched the pulse width modulation signal back and forth between the two motor leads at 2 kHz, so that half the time the AD8210 was seeing a clean current signal and half the time a 30 V 100 kHz common mode signal sitting on top of it. Alternating the PWM signal between motor leads was successful in averaging out the discontinuity at zero when the motor was not spinning, but still left odd variations in current readings with zero current and a spinning motor. The AD8210 is rated to survive positive voltages of up to 68 V, but negative voltages of only -5 V. The small size of the allowable negative common-mode voltage appeared to leave the AD8210 vulnerable to certain types of voltage surges, and we had to replace two of them after using an unusual battery combination.

**Battery voltage sensor:** A resistor network that drops the battery voltage to a reasonable value, which is then buffered by an instrumentation amplifier (AD623) and sent to the ADC. The input voltage range for the AD623 was not really sufficient (the data sheet mentions this limitation in a footnote).

**Inertial Measurement Unit (IMU):** In the first version of Ranger's electronics, the custom microcontroller board included 3 axes each of MEMS accelerometers and rate gyros. However, the rate gyro output seemed quite noisy, so we added a MicroStrain Inertia-Link IMU. The angular position readings were close to what we expected from the data sheet, at 2 degrees accuracy for dynamic measurements and 0.5 degrees for static measurements, at about a 200 Hz data rate. Interestingly, the angular rate data from the Inertia-Link also looked noisy, suggesting that the original sensors were reading correctly after all and that the noise was mechanical. Later we replaced the IMU with the newer MicroStrain 3DM-GX3-25, to reduce power consumption and weight, add magnetic compass capability and achieve a higher data rate. In both cases there were some discrepancies when comparing performance with the data sheets. For the Inertia-Link, we planned for a power consumption of 450 mW because the data sheet said 90 mA and 5 to 9 V. The actual power was over 800 mW; 90 mA apparently applied to the highest possible input voltage, and lower input voltages resulted in higher currents. Similarly, the 3DM-GX3-25 advertises a "1 to 1000 Hz" data rate, without specifying that the top rate applies only to the data coming directly from the accelerometers and rate gyros; using the sensor data to estimate the angular position of the device happens much more slowly, at about the same 200 - 250 Hz of the original Inertia-Link. For its 65 km record walk, Ranger used only the rate gyro information from the IMU, and calculated its own estimates of the leg angles. By using only one rate sensor we did get a 1000 Hz data rate, but without the angular position or magnetic heading data.

IMU data was input to the satellite board via an RS-232 connection to the User Interface (UI) board. The Inertia-Link baud rate was 115 kbaud; for the 3DM-GX3-25 it was increased to 460 kbaud.

**Radio control:** For steering, an older (on hand) Airtronics AM controller was used, with two receiver pulse width outputs measured by MCU timer capture pins on the UI satellite board. Pulse width data can then be used for steering or other remote control functions. There are four timer capture pins available on the board, so a four-channel controller could be used instead.

**Foot contact sensors:** The bottom portion of Ranger's feet are designed to flex up slightly on

contact with the ground, relative to the top portion attached to the ankle joint. Foot flex is measured by a Fairchild H21A1 phototransistor optical interrupter (a $2 device). Although often used as a simple on-off switch for, say, paper handling, the phototransistor versions of these devices make sensitive displacement sensors, with full scale output in fractions of a millimeter. The foot design doesn't give numerical force measurements, since its sensitivity varies with the position of the ground contact along the length of the foot, but as a ground contact sensor it has proven to be highly reliable. The output from the optical sensor goes through the axis of the ankle joint shaft and then out through a hole and up the carbon fiber leg tube. To reduce the risk of wire breakage, we used a highly flex-resistant cable, and spiraled it outward around the ankle shaft to minimize the amount of bending at any point. So far none of these cables have suffered a failure. At the satellite board, the optical sensor is powered and biased as needed, and the output is read by one of the ADC channels.

**Board temperature sensor:** A TI LM26CIM5-XHA temperature IC is mounted close to the MOSFETs on the motor controller, and set to shut down the motor driver if the board temperature exceeds 100 C. It also provides an analog temperature output voltage which goes to one of the internal ADC inputs on the LPC2194/01. We have not seen any board shutdowns due to over-temperature, and did not monitor the temperature readings in software.

**Bluetooth (and direct wire) data logging interface:** A Roving Networks RN21 Bluetooth module is soldered to the CAN router board, and connected by a high-speed serial port to the LPC3250 main processor board. The Bluetooth link (up to 230 kbaud) allows data logging and adjustment of parameters from a remote computer, generally a laptop running Windows. A hardwired connector to a second high-speed serial port (up to 921 kbaud) allows a higher data rate for tethered and bench testing. On the laptop end, a second RN21 module was installed in a custom box along with an FTDI EVAL232R (FT232RL evaluation board) RS232 to USB converter. (The external USB to RS-232 converter is required because Windows internal serial ports, when available on a laptop, do not seem to support baud rates above 115 kbaud.)

**USB and Ethernet ports:** The LPC3250 "main brain" module supports 12 Mbps USB and 10 Mbps Ethernet, so the necessary connectors and support circuitry were added to the CAN router board for possible future use (but have not been used on Ranger yet).

**Two-line LCD display:** A CFAH0802ZYYHJP display (2 lines of 8 characters each) from CrystalFontz is mounted on the UI satellite board and driven through a parallel port connected to MCU digital input-output lines.

**Piezo buzzer:** The buzzer/speaker was driven by a PWM output from the MCU on the UI board, and programmed to beep when errors occurred. Some common errors received their own distinctive alert sounds, and one enterprising student also programmed it to play songs on command, including Cornell's theme "High above Cayuga's Waters" which played once per kilometer during the 65 kilometer walk.

**Keypad:** A six-button keypad, read by the MCU on the user interface board, allows Ranger to be set into various states, including for example "walk," "calibrate," and "standby."

**Color LEDs:** The UI board also has six tri-color LEDs powered by a Maxim MAX6966 LED driver IC, and controlled via the serial peripheral interface (SPI) on the user interface board

MCU. They can be lit up in all different colors, and were used as an easy way to identify Ranger walk controller states, to detect errors from a distance, and to monitor the steering control actions.

### 3.9 Motor controller design

Ranger's motor controllers went through three major changes, roughly corresponding to its 1 km, 9 km, and 65 km walks. These changes led to substantial energy savings. Initially we used the ST VNH2SP30-E motor driver IC, which was intended for things like power windows in cars. It was small, easy to use, and had a high current capability (30 A), but was limited to 20 kHz switching frequency, 16 V maximum, and only switched one end of each half-bridge.

After the 1 km walk we noted that with a higher battery voltage we could get better performance from the nominally 12 V motors. We also noticed that our power losses in the motor windings, due to the ripple current at 20 kHz, were very high. Several motors burned out at average currents that were within the rated maximum, but which created excessive heat buildup due to the high RMS current values resulting from ripple current components. High ripple currents indicate that the pulse-width modulation (PWM) frequency is not high enough for the motor. The motor-controller system has a characteristic inductance and resistance, mostly from the rotor windings, but with small contributions from the controller switch transistors and connecting wires. The motor and controller can be modeled as a series R-L circuit, with a time constant of $L/R$; given some initial value and no external applied voltages, the motor current will decay to zero at time constant $L/R$. One simplistic measure of whether the PWM frequency is high enough is by comparing the time constant $L/R$ to the period of the PWM signal. If the PWM is fast enough, the motor current will not have time to decay appreciably before the next PWM pulse comes along, and the ripple current will be low.

The Faulhaber 2657W012CR is rated for 0.7 Ω terminal resistance ($R$) and 95 $\mu$H winding inductance ($L$). The motor driver "on" resistance is very low, as is the wire resistance. If the loop resistance is assumed to be 0.8 Ω in total (allowing 0.1 Ω for the wires and controller switches), the resulting time constant is $L/R = 120$ $\mu$s, a bit more than twice the 50 $\mu$s PWM period. The ratio of $L/R$ to PWM period should be much higher than this anyway ($L/R$ greater than five times the PWM period would be a good goal), but two additional factors made it much worse. First, the Faulhaber data sheet neglected to add in the effect of the carbon brushes, which according to our later measurements increase the resistance of the total loop to about 1.3 Ω, and introduce a voltage drop of about 0.7 V. Secondly, the VNH2SP30-E only switches the low-side switches at the PWM frequency, leading to a diode voltage drop (about 0.7 V) at the high-side switches. These two voltage drops have the effect of a nonlinear resistor in series with the motor winding, and reduce the $L/R$ time constant inversely with current. For example, at the 3.1 A maximum sustained current rating for the motor, the equivalent added resistance is about $V/I = 1.4/3.1$ Ω $= 0.45$ Ω. As a result, the added voltage drop cuts $L/R$ to only $95/1.75 = 54$ $\mu$s, about equal to the PWM period. For lower currents the effective resistance becomes much higher, and the $L/R$ time even smaller compared to the PWM period, thus guaranteeing high ripple currents with their associated excess heat generation and wasted power.

For the second version of the motor controller electronics, high voltage high-speed drivers were used (Micrel MIC4102) with discrete fast-switching MOSFETs (Toshiba TPCA8014-H). These could use PWM frequencies up to 200 kHz. We used 100 kHz to give a ratio of $L/R$ to PWM period of about five to one (so we thought, see below) at higher winding currents. Higher PWM

frequencies consume more power in driving the MOSFET gate and output capacitances, and reduce the resolution of the PWM duty cycle. (Because of the way the PWM timers work, the duty cycle resolution is at best the PWM frequency divided by the timer clock frequency, here equal to 60 MHz. The resulting resolution equals PWM frequency/60 MHz = 0.17% .

At a 100 kHz PWM frequency the motors did work more efficiently, and this frequency was used for the 9 km record walk. However, upon closer investigation and modeling, the real motor power consumption still did not match what we expected from the simulation. Finally we found that the winding inductance was simply not as high as expected from the data sheet; the motor current waveforms could only be explained by much lower inductance values. Measured on an Agilent 4284A LCR meter, the measured motor inductance dropped from 94 $\mu$H at 1 kHz to 35 $\mu$H at 20 kHz, 20 $\mu$H at 100 kHz, and 16 $\mu$H at 200 kHz. Thus, the loss in inductance largely cancelled out much of the efficiency gains we had hoped for by increasing the PWM frequency. As a check on our equipment and for comparison, we also measured several brands of commercial inductors, which did not show such an inductance drop over our 1 to 1000 kHz test frequency range. We also tested the windings of a Maxon 311536 brushless motor; the Maxon motor inductance was rated at 31 $\mu$H, but in our measurements it dropped to 25 $\mu$H at 100 kHz and 21 $\mu$H at 200 kHz. So inductance drops at higher frequencies are not limited to the Faulhaber motor. We are not clear on the mechanism causing the drop, but suspect parasitic capacitances in the motor winding structure. A consequence of the declining inductance is that increasing the PWM frequency is not a productive way to improve motor efficiency even if the controller can operate efficiently at high PWM frequencies. One common approach to solve this problem and improve efficiency when driving low-inductance motor windings is to add external inductors. For Ranger, we added, in series with each motor, a 47 $\mu$H Vishay/Dale IHLP6767GZER470M11 miniature surface mount inductor with an 8.7-amp maximum current and 0.04 $\Omega$ series resistance. The added inductance brought the excess winding power losses down dramatically, brought the measured power consumption much closer in line with the simulated values, and significantly improved the energy efficiency between the 9 km and 65 km walks.

The motor driver ICs receive a pair of electrically isolated (Analog Devices ADUM1210) PWM signals generated by the PWM outputs of the LPC2194/01 MCU. They also receive a pair of isolated low-side-switch enable signals from MCU digital output lines. These allow the low-side MOSFETs to be shut off at low output currents to save power, if desired (To keep the controller simple, the low-side MOSFETs were always enabled on Ranger). For safety, a watchdog timer monitors another of the MCU digital outputs, and shuts off the power to the motor driver circuitry within about 10 ms if it doesn't see a state transition on the line in that period of time. The digital output line for the watchdog timer, for correct operation, must be toggled by a software function called by the motor control software. Thus, if the motor control software crashes or otherwise stops running, the output line becomes stuck in a fixed state, the timer runs out, and the motors are shut off. Without such a watchdog feature, the motors can turn on unpredictably during programming and debugging, resulting in damage to the hardware and possibly to nearby humans.

### 3.10 Batteries

Ranger uses lithium-ion batteries based on the 18650 cell used in many laptop batteries. The initial version of Ranger used two 11.1 V, 53-watt-hour-nominal batteries in parallel, for a total of about 106 watt hours. However, the voltage of a lithium-ion battery drops substantially as it discharges, from (for an 11.1-volt-nominal battery) a maximum of 12.6 V down to a minimum cut-off voltage

of 9 V. To maintain consistency in the walk controller actions regardless of battery voltage, the motor controllers were compensated to adjust for the drop. The result of the compensation was that the motors always performed as if they were operating at the 9 V minimum voltage, which was not sufficient to give full performance from the 12 V rated motors.

Therefore, with its new higher-voltage motor control boards, Ranger also was fitted with higher-voltage batteries. Initially there were four of these 25.9 V, 67-watt-hour-nominal batteries, totaling about 270 watt hours. Three were strapped to the front of each of the three leg boxes and the fourth to the back of the left leg box; all were positioned to balance the dynamics of the inner vs. outer leg pairs. This battery configuration was used for the 9 km record walk. For the 65 km distance walk, seven of the 25.9 V batteries were used, for a nominal 470 watt-hours of total charge. The battery on the back was removed, and a pair of batteries was added to the right and left outer hips, once again adjusting the positions carefully to match the dynamics of the inner and outer leg pairs.

All of the batteries were from Powerizer/Batteryspace.com, and came installed with overcharge, undercharge, overcurrent, and (for the 25.9 V batteries) cell charge balance circuitry. The 11.1 V batteries could be charged in parallel, but Powerizer technical support recommended against parallel charging for the 25.9 V batteries, apparently because it could disrupt the operation of the charge balance circuits. Thus the 25.9 V batteries had to be charged individually, checked by voltmeter for full and equal charge voltage, and only then attached to the robot and connected in parallel for operation, a procedure that seems risky and was certainly cumbersome.

### 3.11 Network hardware

**Selection of a network (CAN and other alternative networks).** Ranger uses CAN bus for communications between the various satellite boards. CAN was developed for automobile network applications, and is also used in various industrial applications. It is designed for reliable, real-time sensor and control applications in an electrically noisy environment, which also makes it a good choice for a robot. CAN is a peer-to-peer network, so satellite boards can either communicate directly with each other, or transmit all data up to the main processor and receive all commands from it. Compared to Ethernet (and its real-time relative, EtherCAT), CAN is slower, but uses less power, less board space, and is easier to implement, making CAN more suitable for very small, low-power satellite boards. USB is not difficult to implement in hardware, but the network protocol is much more complex and less flexible, though also faster than CAN. CAN has the additional advantage of allowing many nodes on a single cable, thus further reducing the hardware requirements (USB and most Ethernet uses point-to-point wiring; active hubs are needed to connect more than two devices.) Robots in other labs have used networks based on I2C, SPI, RS-485, etc., but these suffer from a relative lack of integration of their network protocol into the controller peripheral on the microcontroller IC. CAN controllers are highly integrated, and include such useful features as CRC error detection, bus arbitration, frame buffers, etc., making transmission of a data frame as easy for the programmer as putting it in the controller transmit buffer; similarly, a received frame can simply be read from the receive buffer when it arrives.

**CAN bus overview.** At the hardware level, a CAN bus is made up of a single length of cable, with one pair of conductors used as a differential pair. The differential nature of the signal helps greatly with noise rejection, because noise signals common to both conductors are subtracted out at the receivers. On Ranger a six-conductor ribbon cable is used, allowing enough space for two

CAN buses and one pair of 5.5 V power supply wires. On Ranger these are unshielded and not twisted, but shielded and twisted-pair versions are also available if more noise rejection is needed. So far, errors in CAN packets have not been a problem, though there was some evidence that the CAN bus was causing interference in an encoder cable. CAN transceivers (nodes) can be attached anywhere along the length of the CAN bus, but it is important that the length of the drop wire between the transceiver and the bus be short. Otherwise, the variation in cable impedance at the tap points can cause the signal to reflect back and possibly interfere with the main signal. On the Ranger satellite boards, the drop wire length is kept to a minimum by having the CAN ribbon cable visit each satellite board in turn. The AMP Micro-MaTch connectors we used can be crimped onto the CAN bus cable at any point, and then plugged into the board. Thus the drop length includes only the connector pins and the short board traces to the CAN transceiver IC. Adding additional boards to the bus is easy - just crimp on another connector somewhere along the cable. These new connections can't really be removed without damage to the cable, but unused bus taps can be plugged up to prevent shorts. Another benefit of using standard-size ribbon cable is the availability of compatible 3M series 3319 high-flex ribbon cable. Used to span rotating joints in the robot, it is rated for a flex life of 100 million cycles.

**CAN bus speed.** The maximum rated speed for CAN bus is 1 Mbps (bits per second). The speed limitation arises from the timing requirements of the arbitration scheme used to resolve data frame collisions, when two nodes begin transmitting at the same time. In this case both transmit together until, eventually, one node transmits a zero and the other a one. Then the zero is dominant and appears on the bus. A station transmitting a one but seeing a zero on the bus recognizes that it has lost the arbitration and stops transmitting. This arbitration scheme allows efficient and predictable use of the bus, because the higher-priority frame always gets through and no time is lost to arbitration. However, if the transmissions are not well synchronized, the arbitration process will fail.

**Overclocking the CAN bus.** Electric signals travel at about 150 meters in a microsecond (about half the speed of light). For 1 Mbps the maximum cable length to maintain synchronization is about 40 m, approximately the scale of a car wiring harness. However, for a small robot like Ranger the bus cables are only a meter or two long. In conjunction with a fast CAN transceiver IC (we used the Maxim MAX3051), the short bus length opens up potentially higher speeds. Many microcontrollers with CAN controllers built in can use higher clock rates to give higher CAN bus speeds, up to about 6 Mbps in the case of the LPC2194/01. In our tests at 6 Mbps the bus did not work no matter how short the cable, and at 5 Mbps it worked only with very short cables, under a meter. At 4 Mbps, however, the performance was good over a 10-meter cable. (We did discover eventually that the acceptance filter for automatically sorting incoming frames did not work for the LPC2194/01 at such high speeds, but it wasn't clear that it worked properly at the normal 1 Mbps speed either - there is mention in the LPC2194/01 errata sheet of problems related to the acceptance filter.)

**CAN bus data rate.** For Ranger we used four independent CAN buses each running at 3 Mbps, which worked well with the 60 MHz processor clock speed and did not push the timing requirements quite so hard. Note that using the higher CAN speeds requires a self-contained set of boards and transceivers selected and set to operate at the non-standard speed. Commercial CAN-compatible

boards were not used in Ranger, but conceivably could have been if one of the buses had been set to operate at 1 Mbps. Each of our CAN frames is 108 bits long, using the maximum 8-byte data payload. Assuming that we allow a maximum bus loading of 75 percent, that means Ranger can handle at most 0.75 x 4 x 3Mbps / 108 = 83,000 frames per second. At first glance 83,000 sounds like a lot, but we want our walk controller to react to sensory information within a few milliseconds. Thus, many of the sensor values need to be sent at 500 Hz, using up the available bandwidth rather quickly. On Ranger we were able to send all of the data we wanted, but some less critical values were sent at a lower rate. For a more complex robot, we would need to pay more attention to data compression, use more CAN buses, or both.

**CAN bus layout.** The CAN Router/Main Brain Carrier board is connected to all four CAN buses; the other satellite boards are connected to two each, though in practice each board only uses one of them for communication. The CAN router's main job is receiving CAN packets from the four buses, checking its lookup table for the correct destination(s) for each frame (based on its 11-bit data ID), and sending them out again. It is the main timekeeper for the system, and for synchronization sends out a CAN frame with the current elapsed time every millisecond to all the other satellites. It also handles high-speed SPI (serial peripheral interface) communications to and from the top-level LPC3250 processor.

### 3.12 Loading code into the microcontrollers

We used the Keil uVision integrated development environment (IDE) for ARM processors, and programmed the flash memory on the MCUs via their JTAG ports using Keil's ULINK debug adapter box. For the satellites, the normal 20-pin 0.1-inch-pitch JTAG port was replaced with a 10-pin AMP Micro-MaTch connector to save space, and used with a small 10-pin adapter board that plugged into the end of the 20-conductor cable from the ULINK. The LPC3250 module has a 20-pin JTAG port onboard; this port was extended to a 20-pin jack on the back of Ranger for convenience. Programming the satellite boards requires removing the thigh box covers, but reprogramming is seldom needed after the code is running well. In principle the satellite code could be distributed automatically via the CAN bus, but so far we have programmed the satellites directly.

## 4 Software for the robot microcontrollers

Later sections of this appendix describe software used for simulation and controller development. Here we describe the software running on the robot itself. By far, this is dominated by code associated with sensing, communications, and low-level motor control. The logic that controls the walking is a small part of the total code.

### 4.1 Overview

**Main brain vs satellites.** Ranger's software architecture represents a balance between two goals: moving processing tasks down to the satellite boards whenever feasible, to improve response time, free up top-level processor resources, and minimize data traffic; and having a single, convenient block of code to work with while improving the performance of the walk controller (compartmentalizing the code). The resulting division of labor between the top-level processor and the various satellites

is relatively conventional: the satellites handle high-speed motor control loops and sensor data inputs, while the top-level processor runs the main estimation and state machine code.

**Modular design.** Paralleling the modular satellite board electronics, the software for the satellite boards is also mostly modular. Each board includes central code modules which handle functions common to all; a subset of the available peripheral driver modules, selected to support the features needed for a particular board; and several setup files to tie all the code together and integrate it with hardware-specific board features.

**Data table.** Data integration and communication across the networked system of processors starts with a robot-wide master table of variables, each with a unique 16-bit data identifier. The table includes parameters that are fixed at the outset and ones that are measured and calculated, including: motor powers, currents, angles and angular rates; battery current and voltage; estimatied dynamic state variables like absolute let angles; number of steps; estimated distance travelled; and various finite state machine parameter (e.g., nominal control level, stiffness, dampness and gain constants). A publish-subscribe system is used; the value associated with any particular data id is generated and published by a single source somewhere in the system, but multiple data destinations (subscribers) are allowed. These data sources and destinations are included in the table, along with an initial value and network routing information. The table is generated as a text file offline, then parsed by a MATLAB script, which generates dedicated C code files for each of the boards in the robot, and also generates code for the data logging and monitoring program. Keeping data definitions consistent across the network is thus automated. In some cases, changes to the master variable table require that all of the code for the robot be recompiled and reloaded onto the boards. However, by adding spare variable slots in the table for each board, we usually only need to reload the code for one or two boards at each onboard update.

**Data sharing.** At run time, the top-level processor on Ranger maintains a structure containing all of the robot's data. The satellite boards only store locally relevant data. When a variable is updated anywhere in the system, it is added to the local schedule to be transmitted as needed to the boards and processes on the subscriber list.

### 4.2 High Level Control: Hierarchical Concurrent Augmented Finite State Machine (FSM)

Ranger's high-level walking control runs in a hierarchical concurrent finite state machine (FSM) on the top-level processor, written in C++. Each motor in the robot - the hip, two ankles, and one steering motor - is assigned its own state machine. These four state machines run concurrently and independently of each other, but are implicitly synchronized by the robot's global dynamic state (the shared data table). For example, heel-strike data can trigger transitions in several state machines at once. And within a state the shared (augmented) data is used in the control. The hierarchical aspect of the state machine was not used in the present controller design, but it is possible to nest one state machine within another, and thus create a multi-level state machine. For example, a "behavior" state could include both a "stand-still" state, a "start-walking" state and a "walk-forwards" state, and each of these could contain four state machines corresponding to the four motors.

**States and transitions.** Each state machine (one for each motor) is made up of a set of states, and a set of transition conditions between states. Each state includes three actions: an entry action function, which runs upon entry to the state; a main action, which runs a control action as long as the state is active, using the 'augmented' data from the table; and an exit action, which runs upon exit from the state. The entry actions and exit actions enable initialization and cleanup of temporary variables, respectively. The transition conditions (e.g., heel-strike), establish the rules for transition from one state to another.

**Compliant control.** Within a state the instantaneous control action typically involves sending a motor command value $I_{\mathrm{i0}}$ to the motor controller at the target joint. This value, nominally a current, is modified by two additional parameters sent previously (possibly in the state's entry action), a proportional gain $K_{\mathrm{ip}}$ and a derivative gain $K_{\mathrm{iv}}$, such that the mechanical impedance at the joint can be controlled. The desired actual motor current $I_{\mathrm{ia}}$ is calculated at the satellite board as:

$$I_{\mathrm{ia}} = I_{\mathrm{i0}} + K_{\mathrm{ip}}\theta_{\mathrm{ia}} + K_{\mathrm{iv}}\dot{\theta}_{\mathrm{ia}} \tag{3}$$

Here $\theta_{\mathrm{ia}}$ and $\dot{\theta}_{\mathrm{ia}}$ correspond to the motor position and motor velocity of the joint that is being controlled.

### 4.3 Sensor fusion and state estimation

The state estimation module on the main brain tries to determine the dynamic state of the robot. It determines which foot is on the ground and the absolute angle and angular rate of the stance leg. The estimator code also keeps track of step time, step velocity, step length, number of steps, distance traveled, etc.

**Dynamic state and heel-strike timing.** Ranger's control depends on knowledge of the absolute angles of the legs relative to the direction of Earth's gravitational acceleration. Especially important is the height of the swing foot from the floor so as to prevent toe-stubbing and to control the step length. These measurements are not directly available from Ranger's onboard sensors, and must be derived via sensor fusion. The MicroStrain IMU (Inertial Measurement Unit) does have an internal general-purpose sensor fusion algorithm for determination of absolute angle, but the data rate of the rate sensor when used alone is higher, and our sensor fusion, based on ground contact and robot kinematics, is more accurate.

On Ranger we use an algorithm which determines absolute leg angles by fusing the sagittal plane angular velocity data from the IMU with other onboard angle and contact sensor data. At heel-strike the configuration of the robot relative to an assumed flat level floor is fully known from joint angle data; at that point the angular rate integrator is partially reset to reflect the newly-measured robot pose relative to the floor, and integration continues from that point forward through the step to the next heel-strike. The degree of reset per step was selected to give a reasonable balance between errors in measurement at heel-strike, roughness of the floor, and the drift rate of the angular velocity measurement. For Ranger, the integrator reset is 10% per step. That is, just after heel-strike the estimate of angle of the IMU is taken as

$$0.9 \cdot (\text{value from integration from last step}) + 0.1 \cdot (\text{value from joint angles assuming level floor}).$$

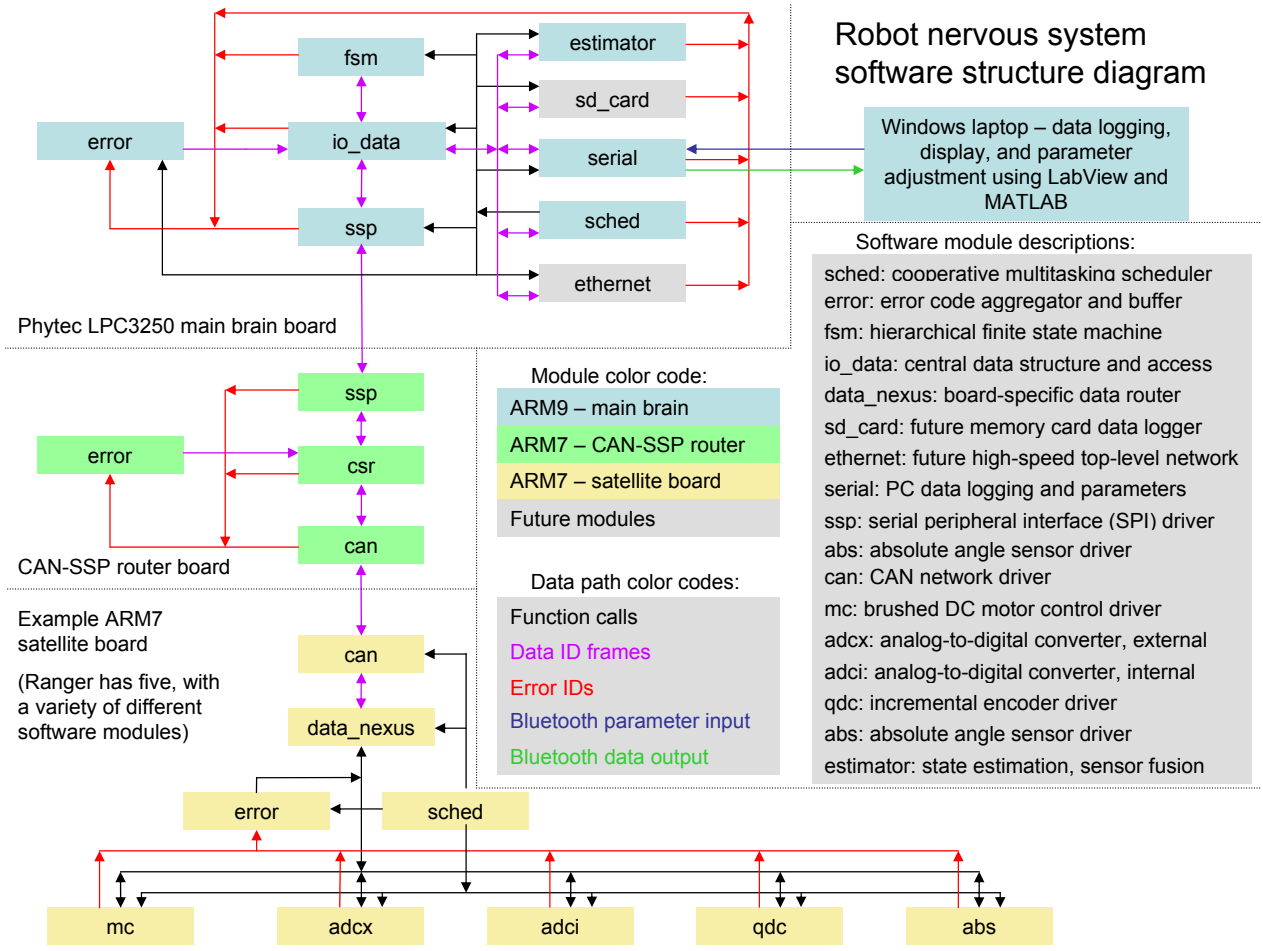


Figure 11: **Overall organization of onboard software.**

Note that the state estimator in its current form is not doing 'model-based-estimation' of the type of common observers or Kalman filters. But the integration reset at each step is a simple kind of model-based estimation. This algorithm, using the angles determined above and the geometry of the kinematic chain knowing the stance foot was on the ground, could measure swing foot height to an accuracy of about 2mm. Our lab floor has step to step variation that is also about 2mm. Hence on the lab floor the swing foot collision could be anticipated with a typical accuracy of about 3-4 mm. As discussed later, accurate prediction of the time of heel-strike based on this dynamical state estimation is enhanced by the foot having a large vertical velocity at heel-strike.

**Finding the time of heel-strike.** The accuracy of the angle estimation (above) depends on the accuracy of the heel-strike detection algorithm; any error in the estimate of the time of heel-strike translates into errors in the estimated absolute leg angles. Heel-strike is detected by ground contact sensors on Ranger's feet. These generate an analog voltage roughly proportional to the force on the heel from the ground. This voltage can have a variety of errors, including high-frequency electrical noise and drift due to temperature, creep, etc. Identifying heel-strike using the unfiltered voltage

and waiting for only a small voltage change could lead to spurious detections. On the other hand, excessive filtering, or waiting for a large change in the force, adds delays to the detection time. Although retrospective detection is fine for integration of angular rates, the controller also has a logical-state transition at heel-strike; delayed detection of heel-strike is problematic for push-off. The heel-strike algorithm on Ranger uses the following algorithm:

1. The swing feet contact sensor voltages are averaged for a period of time during each step while the feet are in the air, well before heel-strike is expected, thus setting up a baseline for comparison and removing long-time sensor drift errors. This zeroing is done while the swing leg is forward of the plane of the hip and the foot has not yet flipped down for heel-strike.

2. Starting from when the foot is flipping down into position for heel-strike, sensor voltages are measured relative to the just-determined baseline. If this baseline voltage for one of the feet is outside of the normal range, that foot's sensor is considered stuck or otherwise in error and ignored for that step.

3. When the contact sensor voltage rises above a predetermined threshold and stays there for at least 2 ms, on either foot, heel-strike is declared and the time is recorded.

The average heel-strike detection delay (time difference between an actual heel-strike and its software detection) is $14 \pm 3$ ms. This delay could be reduced if the threshold were decreased (or the 2 ms limit were decreased). But both of these would reduce the reliability of the detection: a spurious noise or a scuff on debris could cause a false detection of heel-strike.

**Angles and Angular rate estimates.** The angle sensor data from the joint angle sensors is fairly clean (e.g. 0.0008 rad resolution for hip angle, and 0.001 for ankle angles), but the angular rate data is quite noisy. The estimation module filters the local angular rates from the joints and the absolute angular rate from the IMU (which provides a measurement of the absolute angular velocity of the outer leg pair, to which it is attached). A second order Butterworth filter with a cut-off frequency of 10 Hz is used for all of the rate data. A Butterworth filter was chosen because of its maximally flat characteristics in the pass and stop bands. A cut-off of 10 Hz was chosen by looking at the frequencies of the prevalent noise (fluctuations) in the data. The filtering process introduces a phase lag (delay) between the raw and the filtered signal. As always for a filter of given order, the delay and 'smoothing-ability' work against each other. The delay for the current filters is about 20 ms. The filter coefficients were tuned manually by judging delay versus smoothing capability on a test sample of data. Also note that the joint angle rate data received in the main brain has already been filtered in the satellite using a first order filter and hence there is some additional delay of about 5 ms.

### 4.4 Low-level software overview

**Operating system.** The main processor and satellite boards do not run any commercial operating system. Instead they rely on a simple cooperative-multitasking scheduler function. All of the low-level code is written in C, except for a few peripheral drivers written in assembly language. The main processor scheduler runs at 2 ms intervals (500 Hz); the motor controllers run at 0.5 ms intervals (2 kHz); and the other boards run at 1 ms intervals (1 kHz).

**Scheduling.** The order and selection of functions to run in any one of these intervals (schedule slots) is determined by a table. Functions can be placed in one of three categories: 1) Run in every slot. 2) Run in one or more slots (rows) of the table, which is typically ten slots in length. 3) Run occasionally - one function from a list is selected to run during each pass-through of the whole table. These options allow the run rate of a function to be set approximately as desired, from the full control loop rate down to 10 Hz or less. One shortcoming of our system, compared to a commercial real-time operating systems (RTOS), is that all of the functions listed for a slot are required to complete before the end of the interval, putting a strict limit on the amount of processing time available. For example, even if a function on the motor control board only runs once every 100 ms, it still needs to complete in a small fraction of 0.5 ms. Guaranteeing that all the functions complete prior to the end of a slot requires that no function can contain an endless loop or wait for an external event to occur. Instead, each function checks to see if there is anything it can or needs to do, does it quickly, and exits. Execution of the functions in each slot is timed, and approaching or exceeding the allowed slot time generates an error message.

**Interrupts.** High-speed sensor and control input-output (I/O) are handled by interrupt functions at a variety of priority levels, a standard MCU feature which allows normal scheduler operation to be stopped temporarily in order to run a time-critical peripheral function instead. Normal scheduler operation resumes once the interrupt function finishes. As an example, Ranger's ankle motor shaft encoders can send out pulses at over 100 kHz; to avoid miscounting or mistiming these pulses, the appropriate interrupt function must be called immediately each time a new pulse arrives, and thus it receives the highest priority level. CAN frames can arrive at rates of up to 25 kHz per bus, also requiring a high-priority interrupt.

For the top-level processor, the general order of operations during each scheduler slot is:

1. Read in all pending data from the satellites and Bluetooth into the main data structure.
2. Run the state estimator functions.
3. Run the walk controller.
4. Write all command data to the main data structure, for transmission to the satellite boards.

For the satellites it is similar:

1. Run sensor data update functions.
2. Read incoming CAN data.
3. Run the motor controller (for example).
4. Transmit data on CAN bus (sensors, error codes, etc.)

## 4.5 Code module overview

The various boards differ in their application and sometimes their capabilities, but they share core scheduling and network communication functions, and many of the peripheral functions too. These are all encapsulated in reusable code modules, with rules governing their interfaces with other modules, the setup code, and the electronics hardware. Code module software configuration,

data input, and data output are all via a set of public function calls, a fairly standard approach to modular code organization. With the interface to the electronics hardware, we make the connection between the software functions and the details of the electronics as clear and straightforward as possible. Rather than trying to abstract away the details of the hardware interface, our approach does the opposite. For example, MCU peripheral setup registers are referred to by exactly the names given to them in the manufacturer's user manual, for easy reference. Register bits set or cleared are easily identified and looked up if needed.

**Code modules.** Each code module is comprised of a set of functions, some for in-module use only (private) and some intended to be called by other modules (public). A generic set of such functions would include:

Software initialization functions, to be run once at MCU startup.

Input functions, to write data to the module

Output functions, to read data from the module

Update functions, running at predefined intervals. This is where the main processing happens, if needed.

Interrupt function, for peripheral modules with tight timing requirements.

Error logging functions.

In addition, code modules working directly with hardware need this hardware to be initialized properly. Each module has a section at the top with example code for setup of the interrupts and other hardware items it needs. Although that example code doesn't run, it can be copied and pasted as needed into board-specific hardware and interrupt setup functions. These functions take care of the setup for all of the hardware on the board, and run at startup along with the software initialization. By keeping the setup all in one place, it becomes much easier to coordinate the hardware resources between the various code modules, and to avoid conflicts.

### 4.6 Quadrature decoder(QDC) module

Incremental encoders are displacement sensors that output a sequence of electrical pulses in proportion to their relative motion. For example, the Ranger ankle motor encoders send out 512 pulses per revolution of the motor shaft, or approximately one per $360/512 = 0.7$ degrees. A single stream of pulses can tell us how far and how fast the shaft has moved relative to its starting position, but not its direction or absolute position. For incremental encoders, the direction problem is solved by adding a second channel of pulses, out of phase with the first by 90 degrees but otherwise identical (this is called quadrature encoding). When the phase of output A leads that of output B, the shaft is spinning one way, and when the phase of A lags B it is spinning the other way. The QDC module is set up to interrupt when (for example) channel A undergoes a transition in either direction. Based on the state of channel B at the time of the transition, the code can determine the direction of the shaft and thus whether to increment or decrement the position counter. In addition, it can estimate the speed of the shaft by timing the interval between the transitions. Note that the QDC code counts each cycle of channel A twice, once on the rising edge and once on the falling edge,

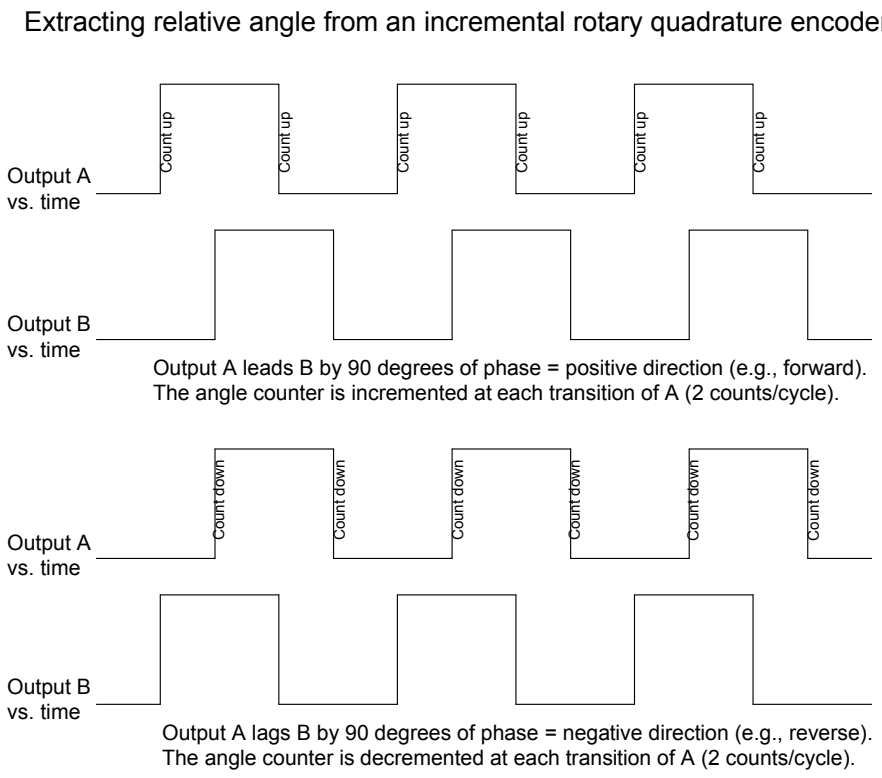
Extracting relative angle from an incremental rotary quadrature encoder
Count up
Count up
Count up
Count up
Count up
Count up
Output A
vs. time
Output B
vs. time
Output A leads B by 90 degrees of phase = positive direction (e.g., forward).
The angle counter is incremented at each transition of A (2 counts/cycle).
Count down
Count down
Count down
Count down
Count down
Count down
Output A
vs. time
Output B
vs. time
Output A lags B by 90 degrees of phase = negative direction (e.g., reverse).
The angle counter is decremented at each transition of A (2 counts/cycle).

Figure 12: **Extracting relative angle from an incremental rotary quadrature encoder.**

giving 2 x 512 or 1024 counts per revolution (2X counting). By also counting the rising and falling edges of channel B this could be doubled to 2048 counts per revolution and four times the base resolution (4X counting). However, the code to do 4X counting requires more processing per count, and has twice as many counts to process as does the 2X counting for a given motor speed; processor speed is an issue.

**Encoder decoding code.** The encoders generate pulses at a high rate when the motors are spinning fast. For example, when an ankle motor is spinning at its full speed, 6300 RPM, the driver software receives pulses at a rate of 6300 * 512 / 60 = 54 kHz, and the interrupt must fire twice per pulse, at 108 kHz. The processor is operating at a clock speed of 60 MHz, so there are (at most, if the processor did nothing else at all) 60 MHz/108 kHz = 558 computation cycles per call to the encoder interrupt function available if it is to keep up with the motor and not lose count. In fact, early versions of our driver did slowly lose count, causing Ranger to fall after a few hundred steps. As a result, code was added to reconcile the ankle motor encoder readings gradually with the respective joint angle sensor readings, preventing the error from growing. In addition, the drivers were rewritten in ARM7 assembly code, reducing the clock cycle count per interrupt function call from about 500 down to about 150. This code optimization lets the motors run at full speed without interfering as much with other running code. In fact, the encoder zeroes no longer seemed to drift with time after the hand coding. Nonetheless, the code to reconcile the incremental-based and absolute angle sensor readings was left in place, just in case.

**Measuring motor velocity.** The encoder interrupt function, besides incrementing or decrementing the position count with each pulse, also sends back the number of clock pulses since the previous encoder pulse. Knowing the pulse period gives us a way to estimate the angular velocity of the motor in addition to its position, but there are some difficulties. First, the encoder pulse widths are not uniform, but have large variations pulse-to-pulse. Second, the number of pulses per second (and thus the data rate) drops as the RPM drops, and becomes zero when the motor is stopped. We wrote a velocity-estimation function that keeps an average of the most recent pulse time values and uses these to estimate the motor velocity. If the pulses stop (very slow or stopped shaft), a second timer keeps track of the time since the most recently received pulse, allowing the function to generate an ongoing estimate of the shaft speed even as it comes to a complete stop.

## 4.7 Motor controller module

**Current control.** At the satellite level, motor control code is built as two nested control loops. The inner loop controls the motor current (a common practice in motor controllers), based on the requested current from the outer loop and feedback from the sensors. At present on Ranger the current controller uses a proportional-integral (PI) control loop, with the emphasis on the integral portion to minimize overshoot. A controller incorporating an additional open-loop (model based) term based on motor speed and battery voltage was demonstrated to give faster response times, but is not yet in use.

**Compliant control.** The outer of the two control loops is a motor compliance controller, using three different control parameters from the main processor to establish a motor current to send to the current controller. The motor current target value is generated as a linear combination of three input parameters: A current offset, a coefficient multiplied by the position, and another coefficient multiplied by the velocity (See Eqn. 3). Compliant control allows the motor controller to behave as a current/torque controller, a position controller, a velocity controller, or any desired combination of the three, through appropriate selection of the three control parameters.

**Coupled motor limits.** Ranger's motors are rated for a 6300 RPM no-load speed at 12 V; i.e., at 12 V and 6300 RPM they generate no output torque. However, with a higher battery voltage, they can generate substantial torque at this speed. A voltage above 12 applied with no load, though, will make the motors spin too fast, and when applied at low speeds may overheat the motors. Thus, separate torque and speed protection functions are required, one keeping the current within safe bounds regardless of the input voltage or speed, and another limiting the speed to a safe value. Moreover, it may be desirable to temporarily drive a motor at currents above the maximum sustained value; short bursts of high current can be used safely if the rotor temperature is not allowed to exceed safe values, and simultaneously, if the peak torque is limited to values that will not cause immediate mechanical damage to the motor or gearhead. For Ranger, therefore, the current limiter has two parts. One incorporates a simple thermal model of the motor armature, based on the rotor thermal time constant and allowed maximum sustained current as given in the data sheet. The measured motor current passes through an exponential averaging filter with a time constant set equal to that of the rotor, and the resulting average is used as the basis of comparison for the thermal limiter. A second limiter looks at the instantaneous (or very short time constant) value of the current, and limits it to give no more than the maximum intermittent torque specification listed on the gearhead data sheet.

## 4.8 Central data module (IO Data)

System-wide parameter and variable data is maintained in a data structure on the top-level ARM9 board. Each variable or parameter to be included is assigned a unique 16-bit data ID for reference. The set of possible data IDs includes but is not equal to the set of CAN network IDs, because the standard CAN ID used in Ranger is limited to 11 bits (actually only the first 2032 of the 2048 possible values in 11 bits, due to CAN implementation details). Higher-value data IDs can be used in the top-level processor, but cannot be transmitted over the CAN network. Each data point in the structure includes a 32-bit data value and a 32-bit time stamp value measured from turn-on time (limiting our walks to 50 days duration). Access to the data structure is through a variety of set- and get-value functions. While the time stamp data type is limited to 32-bit unsigned integer, a union substructure allows the data values to be accessed as any of several data types, so far including single-precision floating point and 32-bit signed and unsigned integers. An additional optional field allows the module to keep track of whether or not subscribing processes have received the latest changes to a data point. This field makes it possible, for example, to limit the data that goes over Bluetooth or CAN to just the data values that have changed, instead of wasting bandwidth by sending a steady stream of identical values. Rather like e-mail, each data value can be marked as read or unread by each subscribing (local) process. There are plenty of additional features that could be added to Ranger's data and network protocol, including the addition of a guaranteed-delivery protocol. One such was largely written, but the data loss rate on the network has been so low that setting this protocol up on Ranger was not a priority. Data types longer than 32 bits would also be good to have, but a large majority of the data values on Ranger are 32-bit single-precision floating point numbers.

## 4.9 Error handling modules

If we want to have clear, reliable data from testing a robot controller, we need to know that the robot is actually running the control code we think it is. It is useless to run tests if the robot is, unbeknownst to us, actually in an error state and thus effectively running a different controller. The error handling modules address this issue. As with the system-wide data IDs described above, the robot also has a system-wide list of 16-bit error codes; these also have a MATLAB script to generate consistent C code for all the boards in the system. When a function in one of the code modules detects an error, it calls the error-occurred function from the error module with the unique error ID number for that condition. The error module maintains a buffer of error messages that been reported so far. Upon receiving another error message, it searches the buffer to see if this is a new error, or a recurrence of an earlier error condition. If new, the error is added to the buffer; if a repeat error, the matching 10-bit error frequency field is incremented. Periodically (as per the scheduler table) an error code is popped off the buffer and transmitted over the CAN bus, receiving as it goes a 6-bit board identifier field and a 32-bit time stamp. The resulting error frame travels over the CAN network to the top-level processor, and then to our data logging and monitoring program (RoboDAQ, described below) on a laptop. Error conditions at the top-level processor are handled similarly, though of course they need not go through the CAN network. RoboDAQ maintains a list of all the detected error conditions, their descriptions, and their total number of occurrences on its error tab, sorted by most recent. It also logs them to the hard disk, as with other received data, and displays the most recent error description and a (virtual) red light on the main data display page.

### 4.10 Foot contact sensor code

This is the low level code used fed to the state estimator for foot contact. Each of Ranger's feet has an integral ground contact sensor based on optical detection of foot structure deflection (described in more detail in the hardware section). The output of the optical sensor is an analog signal which is converted to a voltage measurement in the analog-to-digital converter (ADC). Each 0.5 ms, a new measurement is compared to two threshhold values; when 10 consecutive measurements exceed the upper threshold, the ground contact signal for that foot goes high (true); when the measurements drop below the lower threshhold for ten consecutive times, the foot's ground contact signal goes low (false). This hysteresis and debouncing procedure is intended to keep sensing fast, while minimizing the chance of a false heel-strike detection due to mechanical or electrical noise. At the top-level processor walk controller, the foot-contact signals coming from each pair of feet are logically read such that foot contact on either foot will be interpreted as a heel-strike for the purposes of the walk-control state machines.

### 4.11 RoboDAQ data acquisition program and Bluetooth communications module

This LabVIEW-based software running on a Windows PC communicates data to and from the robot through a high-speed serial port, usually an external FTDI-brand USB to serial converter. From there the data either goes directly to a serial port on the back of Ranger by wire, or through a Bluetooth data link. LabVIEW was selected for ease of programming data acquisition features, and in particular for the ease with which multiple parallel threads of operation can be set up. In this case one thread parses and logs data streaming in from the serial port, another handles data transmission, and a third takes care of the user interface and graph display. The RoboDAQ program includes several useful features for studying the robot's performance:

First, it allows the user to look at graphs of data in real time, while simultaneously logging data to the hard disk. All of the variables in the system, possibly many hundreds, are received and logged periodically, but at a slow 1 Hz rate. However, up to 48 robot variables can be selected for real-time plotting and logging at higher data rates. These requests are relayed to a function on the LPC3250 which puts the desired values into a buffer for transmission to the laptop. By requesting a higher speed, variables can be graphed and logged to a file at up to about 250 Hz over Bluetooth and 500 Hz with the direct RS-232 cable. The actual speed varies depending upon the quality of the wireless link, and how many variables are chosen for higher-speed transmission.

Second, data can be logged to a file on disk for later analysis, generally by a custom MATLAB-based plotting program. One of these MATLAB programs allows synchronization of the data plots with robot video.

Third, users can adjust parameters on the robot while it is walking, using the Bluetooth link. RoboDAQ also allows saving sets of parameters for later use, and loading these saved parameters to the robot. For the walking distance record attempts, parameter adjustment was turned off, and Ranger relied on parameters compiled into its code. Only the data logging features were used.

Fourth, the LabVIEW program reports error conditions from all levels of the Ranger code. Error monitoring is critical, when working with a complex device like Ranger, to make certain that unexpected error conditions are not corrupting the results of the testing.

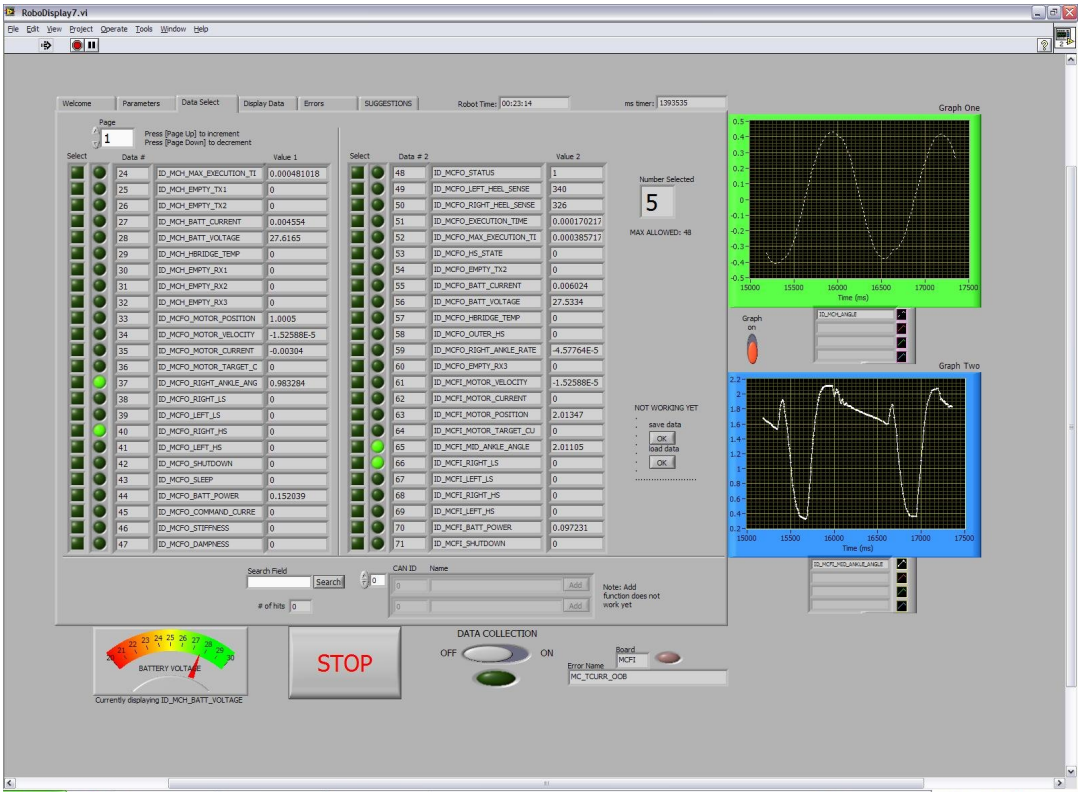


(a) Sensor Variables

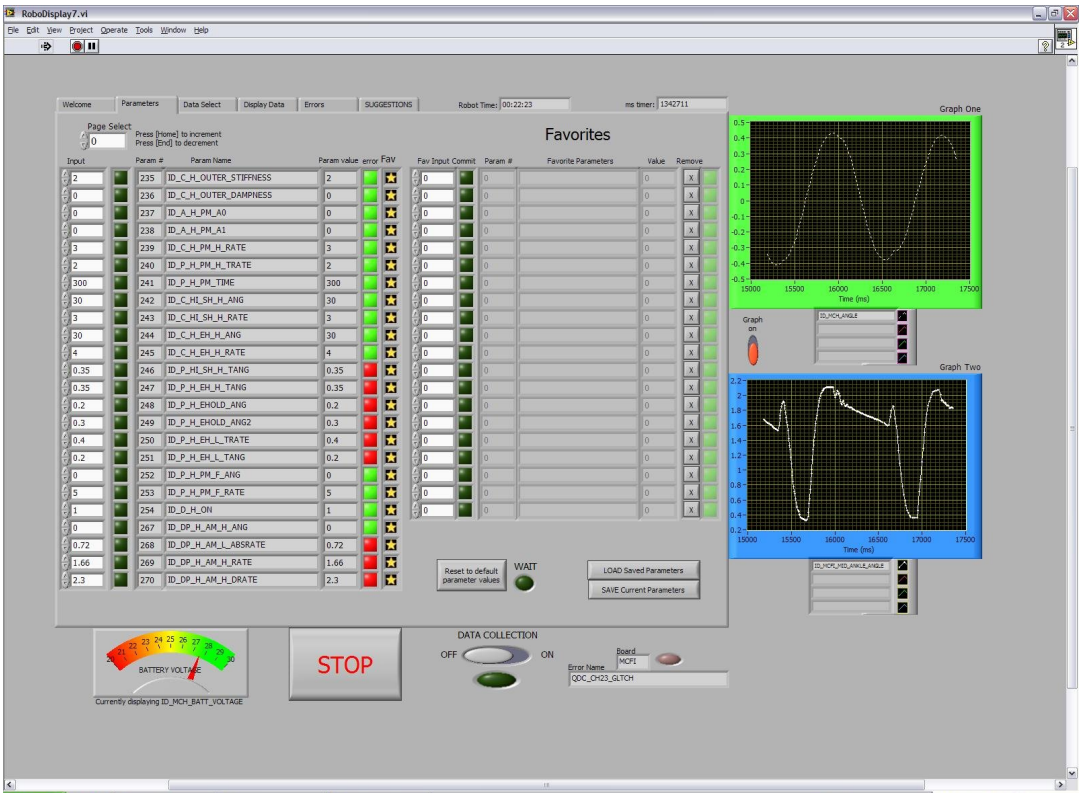


(b) Parameter Variables

Figure 13: **Snapshots of LabVIEW based program on a Windows PC that communicates with the robot.** There is a tab for **(a)** display of sensor variables, and **(b)** display of parameter variables

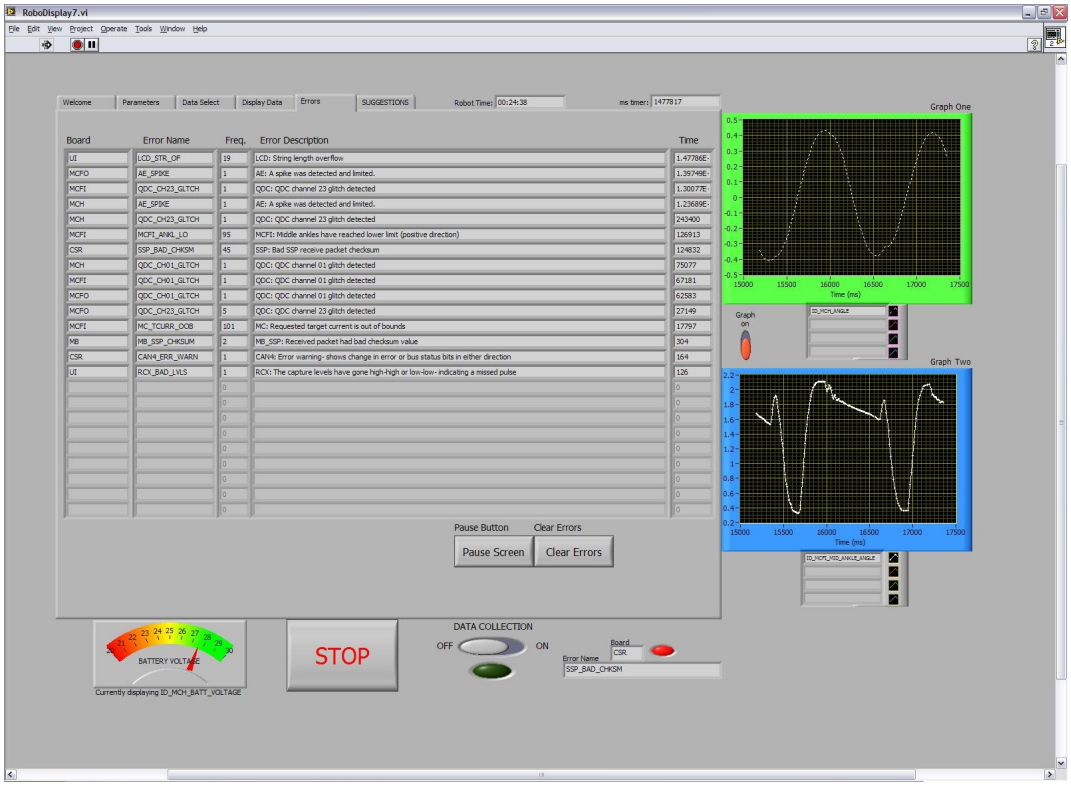


(c) Errors

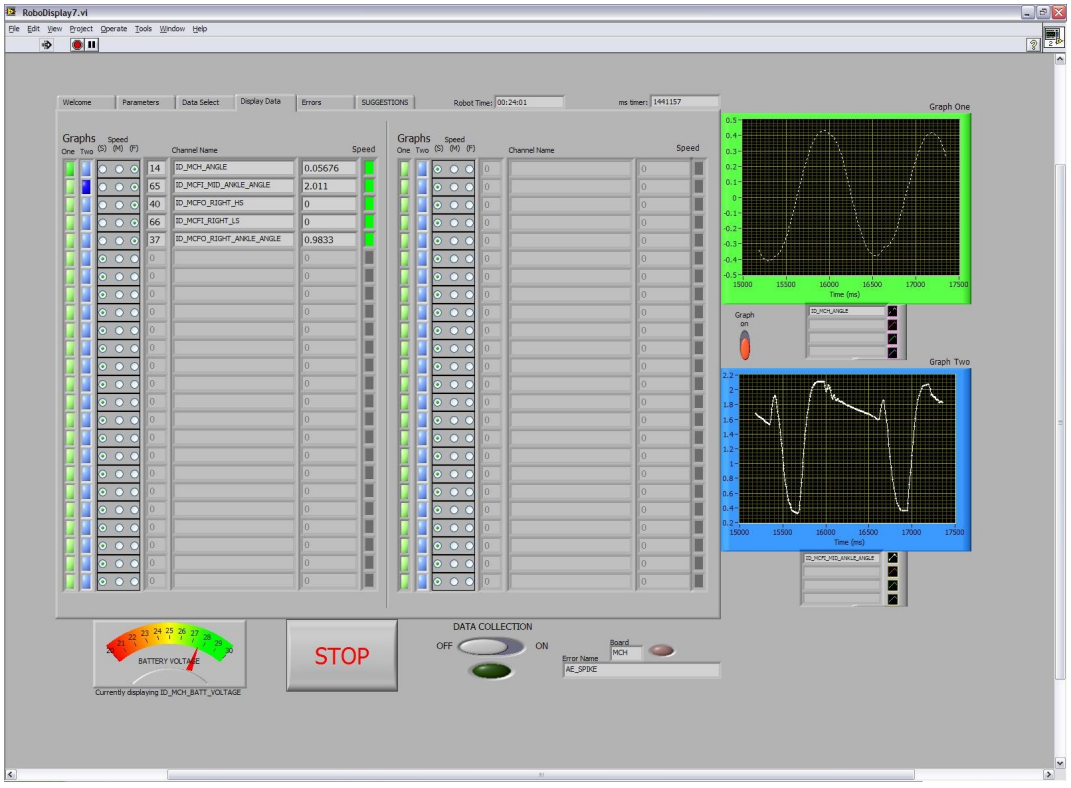


(d) User selects sensor variables for monitoring and data collection.

Figure 14: **Snapshots of LabVIEW based program on a Windows PC that communicates with the robot.** There is a tab for **(c)** display of various errors on the robot and **(d)** tab to monitor and collect data from a user defined set from the sensor variables in (a).

**Bluetooth.** The Bluetooth data module is the RoboDAQ program's counterpart on the robot itself. Running on the top-level ARM9 processor, its job is to look at the data requests from RoboDAQ, updates to the central data structure, and the available transmission bandwidth, and then decide which and how many data points should be put into the serial port buffers and transmitted. For this purpose the code maintains a high-speed data list and a medium-speed data list; all other data is by default slow-speed. For each transmission cycle the Bluetooth module sends all the data in the high-speed list, one data point from the medium-speed list, and one from the slow-speed list. The relative speed of the lists is determined by the lengths of the lists. For example, medium-speed data only gets sent at a higher rate than the slow-speed data if its list is shorter; similarly, if the fast list and the medium list each have only one item, they will be sent at the same fast speed. The slow data is sent whether it has changed since the previous transmission or not; fast and medium-speed data is only transmitted when the values change, to save bandwidth. The transmission rate for a direct-wire serial connection is quite consistent, but the Bluetooth connection speed varies with signal quality, and periodically shuts down for short periods of time, presumably for internal protocol maintenance. These transmission gaps make it difficult to maintain a steady, uninterrupted transmission rate via Bluetooth. A proportional controller was added on the transmit rate, tuned to try to keep a large output buffer half-filled; this helped, but when the Bluetooth link drops out for too long the buffer eventually overflows and leaves a gap in the data log.

## 4.12 Other code modules

Ranger uses about 25 other code modules, ranging from a few small files to read in pushbutton switch values and light up displays and LEDs, to more complex parsing and estimating code for sensors like the IMU and foot contact sensors. See the code itself and associated documentation for more information (available online).

# 5 Equations of motion

The robot model is described in the main paper and shown in figure 1b therein. The dimensions are shown in figure 1 below. Reference frames and angles are shown in figure 15.

**Coordinates.** During double stance we can write the co-ordinates of the hip $x_h, y_h$, using the fixed (Newtonian) coordinate system $xy$ in two ways (see figure 15): 1) using the path $\mathrm{OP_1A_1H}$, and 2) using the path $\mathrm{OP_2A_2H}$. We thus have two kinematic restrictions (constraints) on the joint angles ($q_2, q_3$ and $q_4$), the absolute angle of the lead foot $q_1$ and step length ($x_{p2} - x_{p1}$) during double stance.

$$
\begin{aligned}
x_h &= x_h \\
\implies x_{p1} + l\sin(q_1+q_2) - d\sin(q_1) - rq_1 &= x_{p2} - l\sin(q_3-q_1-q_2) - d\sin(q_1+q_2-q_3-q_4)\ldots \\
&\quad - r(q_1+q_2-q_3-q_4) \qquad (4)\\
y_h &= y_h \\
\implies r - l\cos(q_1+q_2) + d\cos(q_1) &= r - l\cos(q_3-q_1-q_2) + d\cos(q_1+q_2-q_3-q_4) \qquad (5)
\end{aligned}
$$

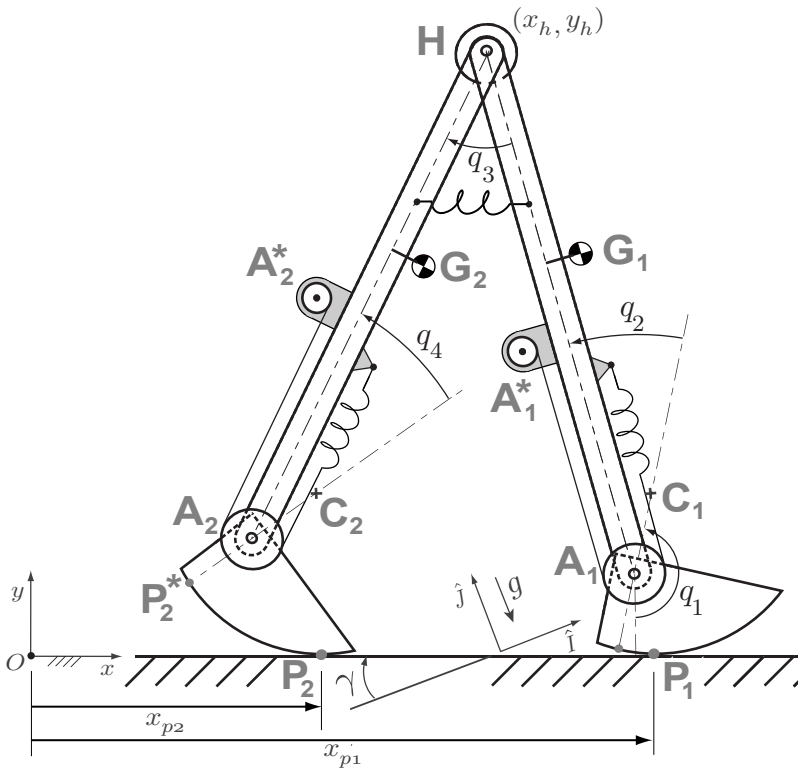


Figure 15: **Robot reference frames and degrees of freedom used in the derivation of the equations of motion.** The absolute angle made by the lead foot on the ground with the vertical is $q_1$. Joint angles are $q_2$, $q_3$ and $q_4$. Hip motor angle is the same as hip joint angle $q_3$. Ankle motor angles associated with the joint $A_1^\star$ is $q_{2m}$ and with joint $A_2^\star$ is $q_{4m}$.

### 5.1 Equations of motion during single and double stance

The governing differential equations are found using the free body diagrams (FBDs) shown in figures 16 and 17, respectively, and using angular momentum balance about judiciously chosen points that eliminate (at least some of) the constraint forces.

**Double stance.** In double stance the robot has 2 kinematic degrees of freedom and two motor degrees of freedom. Releasing one foot from the ground and then constraining it, we think of this rather as 4 kinematic degrees of freedom, 2 motor degrees of freedom and 2 kinematic constraint equations. From figure 16 we use the systems defined by the free body diagrams (a) – (d). We then use angular momentum balance about the points $P_1, A_1, H$ and $A_2$, respectively, to generate

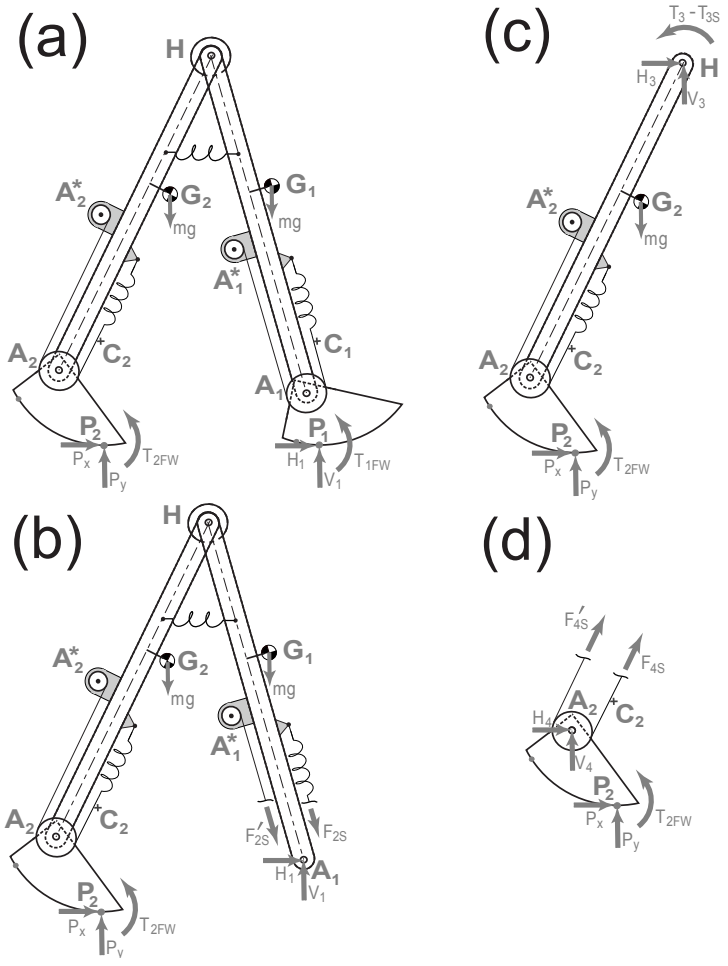


Figure 16: **Free Body Diagrams (FBD) to derive equations for double stance.** We have four free body diagrams. The arrows indicate all of the non-neglected forces and torques acting on each of the four systems.

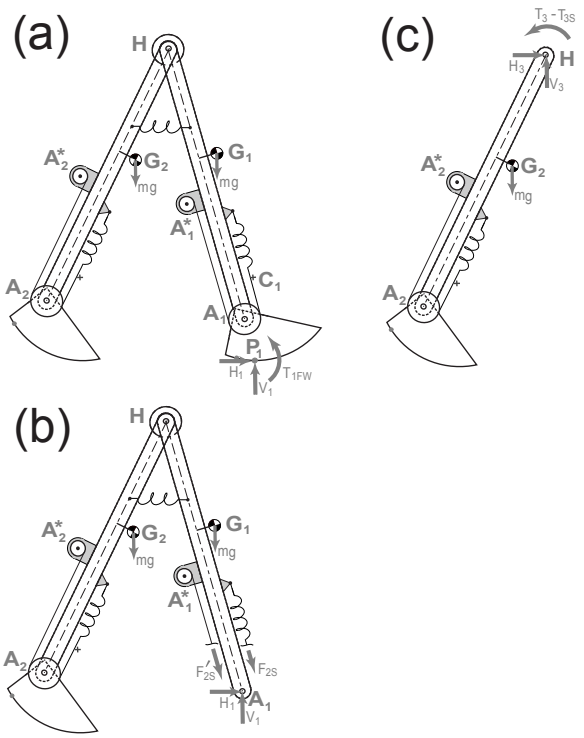


Figure 17: **Free Body Diagrams (FBD) to derive equation for single stance.** We have 3 free body diagrams. Because the feet are massless there is no information in drawing a FBD of the swing foot.

4 equations. The motor degrees of freedom are described with the motor equations from the main paper (equations 7 and 8 below). Two additional constraint equations for double stance are obtained by taking the second derivatives of equations 4 and 5,

$$\dot{\overrightarrow{H}}_{/i} = \overrightarrow{M}_{/i} \qquad \text{where } i = P_1, A_1, H, A_2 \tag{6}$$

$$T_{2S} = G_A(KI_2 - G_A J_m \ddot{q}_{2m}) - T_{fA}(I_2, \dot{q}_{2m}) \tag{7}$$

$$T_{4S} = G_A(KI_4 - G_A J_m \ddot{q}_{4m}) - T_{fA}(I_4, \dot{q}_{4m}) \tag{8}$$

$$\ddot{x}_h = \ddot{x}_h \tag{9}$$

$$\ddot{y}_h = \ddot{y}_h \tag{10}$$

$\overrightarrow{M}_{/i}$ and $\dot{\overrightarrow{H}}_{/i}$ are the sum of external torques and rate of change of angular momentum about the point $i$. These expressions are shown expanded in section 5.3. Note, from the main text the spring torque at joint 2 is, $T_{2S} = F_{2S} r_p = k r_p^2 (q_{2m} - q_2) = k_s (q_{2m} - q_2)$ and spring torque at joint 4 is $T_{4S} = F_{4S} r_p = k r_p^2 (q_{4m} - q_4) = k_s (q_{4m} - q_4)$. $F_{2S}$ and $F_{4S}$ are the spring forces in the springs at joint 2 and 4 respectively, $r_p$ is the radius of the ankle pulley and $k$ is the linear spring constant of the cable joining ankle to the ankle motor. Please see the nomenclature (section 1) for definitions of terms.

Altogether we have 6 differential equations of motion and two differential equations from differ-

entiating the closed-linkage geometric constraint. From these we can solve, at any instant in time, for the angular accelerations of each robot part, both ankle motors, and the ground contact force at one foot.

**Single stance.** The swing-foot is airborne for the single stance phase. Because we neglect the masses of the feet, the swing foot and swing motor have the same motions and the swing foot does not have independent motion. This eliminates one degree of freedom from the single stance phase. Thus, in the single stance phase we have five degrees of freedom (stance ankle angle, hip angle, two motor angles and the stance foot angle). Altogether we then have three kinematic degrees of freedom, and two motor degrees of freedom.

Using figure 17 we use the systems shown in the free body diagrams (a)–(c). With these we use angular momentum balance about the points $P_1, H$ and $A_1$ respectively, to generate 3 equations of motion. Two more equations come from the two ankle motor equations. Thus, we have the following equations

$$\dot{\overrightarrow{H}}_{/i} = \overrightarrow{M}_{/i} \qquad \text{where } i = P_1, A_1, H \tag{11}$$

$$T_{2S} = G_A(KI_2 - G_A J_m \ddot{q}_{2m}) - T_{fA}(I_2, \dot{q}_{2m}) \tag{12}$$

$$0 = G_A(KI_4 - G_A J_m \ddot{q}_{4m}) - T_{fA}(I_4, \dot{q}_{4m}) \tag{13}$$

$\overrightarrow{M}_{/i}$ and $\dot{\overrightarrow{H}}_{/i}$ are the sums of the external torques and the rate of change of angular momentum about the point $i$. In section 5.3, we give the detailed expansions of these expressions. The equations above make up 5 differential equations for the angles of three body parts (all but the swing foot) and the two motors.

**Matrix form of the governing equations.** Equations 6 to 10 can be re-arranged to get the following equation for double stance,

$$\mathbf{A}_{ds}\mathbf{X}_{ds} = \mathbf{b}_{ds} \tag{14}$$

where the unknown is the $8 \times 1$ vector, $\mathbf{X}_{ds} = [\ddot{q}_1\ ,\ \ddot{q}_2\ ,\ \ddot{q}_{2m}\ ,\ \ddot{q}_3\ ,\ \ddot{q}_4\ ,\ \ddot{q}_{4m}\ ,\ P_x\ ,\ P_y]'$. At a given dynamical state (given angles and rates) the $8 \times 8$ matrix $\mathbf{A}_{ds}$ and the $8 \times 1$ vector $\mathbf{b}_{ds}$ are known.

**Extracting the elements of $\mathbf{A}_{ds}$ and $\mathbf{b}_{ds}$.** There are various ways to find the elements of $\mathbf{A}_{ds}$ and $\mathbf{b}_{ds}$. Here is our somewhat clumsy method. We use symbolic algebra to evaluate the eight equations given in equation 6 to equation 10. Our next goal is to compute symbolic values of the individual elements of the matrices in $\mathbf{A}_{ds}$ and vector $\mathbf{b}_{ds}$. The first element of $\mathbf{b}_{ds}$ i.e. $b_1$ is obtained by setting $\mathbf{X}_{ds} = [0\ ,\ 0\ ,\ 0\ ,\ 0\ ,\ 0\ ,\ 0\ ,\ 0\ ,\ 0]$ and evaluating the first equation in 6. Similarly, we can calculate the other elements of $\mathbf{b}_{ds}$. To get the first row and first column element of $\mathbf{A}_{ds}$ i.e. $A_{11}$, we evaluate first equation in 6 by putting $\mathbf{X}_{ds} = [1\ ,\ 0\ ,\ 0\ ,\ 0\ ,\ 0\ ,\ 0\ ,\ 0\ ,\ 0]$ and from this value subtract out $b_1$. Following a similar procedure it is possible to get every element of the matrix $\mathbf{A}_{ds}$. As pointed out by Manoj Srinivasan (private communication) the MATLAB symbolic command JACOBIAN could have simplified this extraction.

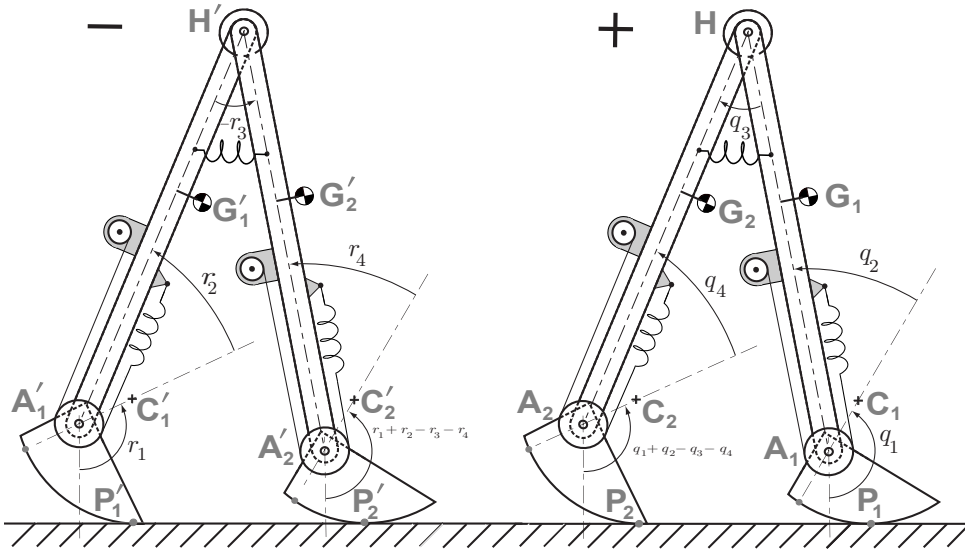


Figure 18: **Angle swap for heel-strike derivations.** The instant just before heel-strike is denoted by $^-$ and the instant just after heel-strike by $^+$. We swap the names of the legs during heel-strike as shown. To simplify notation the angles are named $r_i$ before collision and $q_i$ after collision., where $i$ is the joint number.

**Single stance equations in matrix form.** Equations 11 to 13 can be re-arranged to get the following equation for single stance.

$$\mathbf{A}_{ss}\mathbf{X}_{ss} = \mathbf{b}_{ss} \tag{15}$$

where the unknown is the $6 \times 1$ vector, $\mathbf{X}_{ss} = [\ddot{q}_1\ ,\ \ddot{q}_2\ ,\ \ddot{q}_{2m}\ ,\ \ddot{q}_3\ ,\ \ddot{q}_4\ ,\ \ddot{q}_{4m}\ ]'$, while the $6 \times 6$ matrix $\mathbf{A}_{ss}$ and the $6 \times 1$ vector $\mathbf{b}_{ss}$ are known and can be found in a similar way as found for the double stance equations.

## 5.2 Collisional heel-strike equations

Here we consider the jump (the discontinuity) in angular rates when the swing foot collides with the ground at heel-strike. We consider first the case when this is a transition from a single-stance phase to a double-stance phase.

Figure 18 shows the robot an instant before heel-strike, denoted by $^-$ and an instant after heel-strike, denoted by $^+$. We are interested in finding the angles after heelstrike, i.e.
$[\mathbf{q}] = [q_1\ ,\ q_2\ ,\ q_{2m}\ ,\ q_3\ ,\ q_4\ ,\ q_{4m}]'$ and angular velocities after heel-strike, i.e.
$[\dot{\mathbf{q}}^+] = [\dot{q}_1\ ,\ \dot{q}_2\ ,\ \dot{q}_{2m}\ ,\ \dot{q}_3\ ,\ \dot{q}_4\ ,\ \dot{q}_{4m}]'$.

The angles after heel-strike are found by using figure 18 and swapping the angles to generate equation 16. The velocities of joints after heel-strike are found by conservation of angular momentum. Doing conservation of angular momentum about appropriate points as shown in figure 19 (a)–(d) we get equations 17. We assume that the motors (buffered by the ankle spring) do not participate in the heel-strike. We swap the motor velocities to get equations 18. Finally, two additional constraint equations are generated by taking the first derivatives of equations 4 and 5

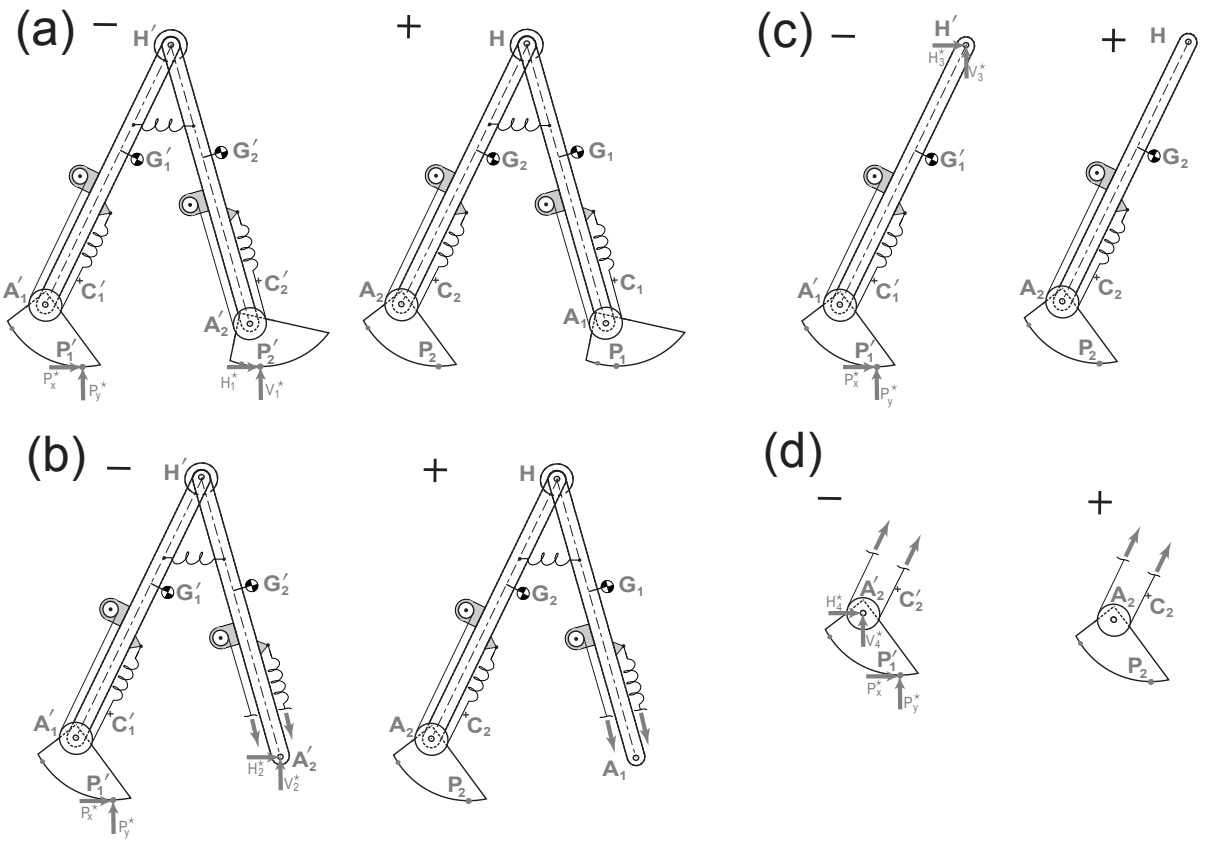


Figure 19: **Free Body Diagrams (FBDs) for heel-strike discontinuity.** The case shown is for the transition to a double-stance phase; there are impulses at both feet. The instant just before and after heel-strike is indicated by $^-$ and $^+$ respectively. The collisional forces are shown in the $^-$ configuration. We have four subsystems corresponding to the four body parts. The ankle motors are buffered by the ankle springs and do not participate in the collision; so the inner and outer ankle motors and rates simply exchange values (ie., keep their values and exchange their names).

to give equation 19.

$$\begin{aligned} q_1 &= r_1 + r_2 - r_3 - r_4 & q_2 &= r_4 \\ q_{2m} &= r_{4m} & q_3 &= -r_3 \\ q_4 &= r_2 & q_{4m} &= r_{2m} \end{aligned} \tag{16}$$

$$\begin{aligned} \overrightarrow{H}^{+}_{/P_1} &= \overrightarrow{H}^{-}_{/P_2'} + \vec{r}_{P_1'/P_2'} \times \vec{P}^{\star} & \overrightarrow{H}^{+}_{/A_1} &= \overrightarrow{H}^{-}_{/A_2'} + \vec{r}_{P_1'/A_2'} \times \vec{P}^{\star} \\ \overrightarrow{H}^{+}_{/H} &= \overrightarrow{H}^{-}_{/H'} + \vec{r}_{P_1'/H} \times \vec{P}^{\star} & \overrightarrow{H}^{+}_{/A_2} &= \overrightarrow{H}^{-}_{/A_1'} + \vec{r}_{P_1'/A_1'} \times \vec{P}^{\star} \end{aligned} \tag{17}$$

$$\dot{q}_{2m} = \dot{r}_{4m} \qquad \dot{q}_{4m} = \dot{r}_{2m} \tag{18}$$

$$\dot{x}_h = \dot{x}_h \qquad \dot{y}_h = \dot{y}_h \tag{19}$$

where $\overrightarrow{H}^{-}_{/i}$ and $\overrightarrow{H}^{+}_{/i}$ are the angular momentum about the point $i$ before and after heel-strke. In section 5.3, we give the detailed expansions of these expressions.

The equations 17 to 19 can be re-arranged to give the following equation for heel-strike phase,

$$\mathbf{A}_{hs}\mathbf{X}_{hs} = \mathbf{b}_{hs} \tag{20}$$

where the unknown is the $8 \times 1$ vector, $\mathbf{X}_{hs} = [\dot{q}_1\ ,\ \dot{q}_2\ ,\ \dot{q}_{2m}\ ,\ \dot{q}_3\ ,\ \dot{q}_4\ ,\ \dot{q}_{4m}\ ,\ P^{\star}_x\ ,\ P^{\star}_y]'$, while the $8 \times 8$ matrix $\mathbf{A}_{hs}$ and the $8 \times 1$ vector $\mathbf{b}_{hs}$ are known and can be found in a similar way as found for the double stance equations.

**Single-stance to single-stance collisional transition.** The equations presented next are for gait sequences that do not have any double stance; i.e., the gait sequence is single stance → heel-strike → single stance. This sequence is not used on ranger, but is used as a special case to benchmark the equations of motion (section 6).

The collisional jump equations for the single stance to single stance transition can be derived by similar means as used for the single stance to double stance transition as described before. We use the same figure 18 for this derivation, but 1) leave off the kinematic constraint that the former stance foot maintain contact; 2) leave off the related constraint impulses (set the impulses on the former stance foot to zero). The jump equations are:

$$q_1 = r_1 + r_2 - r_3 - r_4 \qquad q_3 = -r_3 \tag{21}$$

$$q_{2m} = r_{4m} \qquad q_2 = r_4 \tag{22}$$

$$q_4 = r_2 \qquad q_{4m} = r_{2m} \tag{23}$$

$$\overrightarrow{H}^{+}_{/P_1} = \overrightarrow{H}^{-}_{/P'_2} \tag{24}$$

$$\overrightarrow{H}^{+}_{/A_1} = \overrightarrow{H}^{-}_{/A'_2} \tag{25}$$

$$\overrightarrow{H}^{+}_{/H} = \overrightarrow{H}^{-}_{/H'} \tag{26}$$

$$\dot{q}_4 = \dot{r}_{2m} \qquad \dot{q}_{2m} = \dot{r}_{4m} \qquad \dot{q}_{4m} = \dot{r}_{2m} \tag{27}$$

$\overrightarrow{H}^{-}_{/i}$ and $\overrightarrow{H}^{+}_{/i}$ are the angular momentum about the point $i$ before and after heel-strke. Formulas for these follow.

## 5.3 Full expansion of terms in the single stance, double stance and heel-strike equations

The external moment $\overrightarrow{M}_{/i}$ and rate of change of angular momentum $\dot{\overrightarrow{H}}_{/i}$ about different points in the above equations are given next. We define $\delta_{DS} = 1$ in double stance and $\delta_{DS} = 0$ in single stance in the equations below.

$$\overrightarrow{M}_{/P_1} = \vec{r}_{G_1/P_1} \times m\,\vec{g} + \vec{r}_{G_2/P_1} \times m\,\vec{g} + \vec{r}_{P_2/P_1} \times \delta_{DS}\,\vec{P} + (T_{1FW} + \delta_{DS}T_{2FW})\hat{k}$$

$$
\begin{aligned}
\dot{\overrightarrow{H}}_{/P_1} &= \overrightarrow{r}_{G_1/P_1} \times m\, \overrightarrow{a}_{G_1} + J_\ell\, \overrightarrow{\alpha}_1 + \overrightarrow{r}_{G_2/P_1} \times m\, \overrightarrow{a}_{G_2} + J_\ell\, \overrightarrow{\alpha}_2 \\
\overrightarrow{M}_{/A_1} &= \overrightarrow{r}_{G_1/A_1} \times m\, \overrightarrow{g} + \overrightarrow{r}_{G_2/A_1} \times m\, \overrightarrow{g} + \overrightarrow{r}_{P_2/A_1} \times \delta_{DS}\, \overrightarrow{P} + (T_{2S} + \delta_{DS} T_{2FW})\hat{k} \\
\dot{\overrightarrow{H}}_{/A_1} &= \overrightarrow{r}_{G_1/A_1} \times m\, \overrightarrow{a}_{G_1} + J_\ell\, \overrightarrow{\alpha}_1 + \overrightarrow{r}_{G_2/A_1} \times m\, \overrightarrow{a}_{G_2} + J_\ell\, \overrightarrow{\alpha}_2 \\
\overrightarrow{M}_{/H} &= \overrightarrow{r}_{G_2/H} \times m\, \overrightarrow{g} + \overrightarrow{r}_{P_2/H} \times \delta_{DS}\, \overrightarrow{P} + (T_3 - T_{3S} + \delta_{DS} T_{2FW})\hat{k} \\
\dot{\overrightarrow{H}}_{/H} &= \overrightarrow{r}_{G_2/H} \times m\, \overrightarrow{a}_{G_2} + J_\ell\, \overrightarrow{\alpha}_2 \\
\overrightarrow{M}_{/A_2} &= \overrightarrow{r}_{P_2/A_2} \times \overrightarrow{P} + (T_{4S} + T_{2FW})\hat{k} \qquad \text{Only in double stance} \\
\dot{\overrightarrow{H}}_{/A_2} &= 0 \qquad \text{Only in double stance}
\end{aligned}
$$

The angular momentum $\overrightarrow{H}_{/i}$ about different points in the heel-strike equations are given next.

$$
\begin{aligned}
\overrightarrow{H}^{+}_{/P_1} &= \overrightarrow{r}_{G_1/P_1} \times m\, \overrightarrow{v}_{G_1} + J_\ell\, \overrightarrow{\omega}_1 + \overrightarrow{r}_{G_2/P_1} \times m\, \overrightarrow{v}_{G_2} + J_\ell\, \overrightarrow{\omega}_2 \\
\overrightarrow{H}^{-}_{/P_2'} &= \overrightarrow{r}_{G_1'/P_2'} \times m\, \overrightarrow{v}_{G_2'} + J_\ell\, \overrightarrow{\omega'}_1 + \overrightarrow{r}_{G_2'/P_2'} \times m\, \overrightarrow{v}_{G_2'} + J_\ell\, \overrightarrow{\omega'}_2 \\
\overrightarrow{H}^{+}_{/A_1} &= \overrightarrow{r}_{G_1/A_1} \times m\, \overrightarrow{v}_{A_1} + J_\ell\, \overrightarrow{\omega}_1 + \overrightarrow{r}_{G_2/A_1} \times m\, \overrightarrow{v}_{G_2} + J_\ell\, \overrightarrow{\omega}_2 \\
\overrightarrow{H}^{-}_{/A_2'} &= \overrightarrow{r}_{G_1'/A_2'} \times m\, \overrightarrow{v}_{G_2'} + J_\ell\, \overrightarrow{\omega'}_1 + \overrightarrow{r}_{G_2'/A_2'} \times m\, \overrightarrow{v}_{G_2'} + J_\ell\, \overrightarrow{\omega'}_2 \\
\overrightarrow{H}^{+}_{/H} &= \overrightarrow{r}_{G_2/H} \times m\, \overrightarrow{v}_{G_2} + J_\ell\, \overrightarrow{\omega}_2 \\
\overrightarrow{H}^{-}_{/H'} &= \overrightarrow{r}_{G_1'/H'} \times m\, \overrightarrow{v}_{G_1'} + J_\ell\, \overrightarrow{\omega'}_1 \\
\overrightarrow{H}^{+}_{/A_2} &= \overrightarrow{H}^{-}_{/A_1'} = \overrightarrow{0}
\end{aligned}
$$

# 6 Benchmark tests of the equations of motion

To validate the equations of motion we consider various special cases about which much is known. In particular, we reduced the equations of motion of Ranger to that of a passive walker, introduced a ramp, found stable limit cycles and compared our results with previously published results. Two special passive cases were considered: 1) a rimless wheel [1, 8], and the simplest walker [5, 6]. Without such checks how are we to trust our equations?

## 6.1 Recipe for analyzing passive dynamic walkers

A recipe for analyzing passive dynamic walkers has been presented in detail in Garcia's PhD thesis [6] (see Chapter 2) and in Coleman's PhD thesis [1] (see Chapter 1). Here is a summary.

1. Create a complete mechanical model of the walker. This includes defining model assumptions, defining model parameters, and deriving the equations of motion of the walker. The equations of motion generally involve smooth portions of the walk like single stance phase and discontinuous portions like heel-strike.

2. Define a Poincare map (McGeer's stride function) that maps the state of the system from one step to the next. Here we used the map $f$ from the state of the system $q_n$ just after one heel-strike to the state $q_{n+1}$ just after the next heel-strike. We use numerical root finding (e.g. Newton-Raphson method or bisection method) to find fixed points (roots of the function $f(q) - q$).

3. Find the Jacobian of the stride function at this root using numerical differentiation. If the magnitude of the biggest eigenvalue of the Jacobian is smaller than one than the system is stable, otherwise it is not.

## 6.2 Reduction to a 2-D rimless wheel

Figure 20 (a) shows the 2-D rimless wheel analyzed by Coleman. The mass at the center is $M$, legs have inertia $I_\ell$, legs length is $\ell$ and number of spokes is $n$. The inter-spoke angle $\beta$ is calculated from the number of spokes and is given by $\beta = 2\pi/n$. The ramp slope is $\gamma$. Coleman found that the roots and Jacobian of the stride function depend on the number of spokes $n$, the slope $\gamma$ and the non-dimensional term $\lambda^2 = M\ell^2/(I_\ell + M\ell^2)$. Coleman reports analytical results for $n = 6$, $\gamma = 0.2$ and $\lambda^2 = 2/3$ [2]. We will use these parameters for the benchmark.

**Ranger model reduced to a rimless wheel.** To derive the equations of motion we proceed as follows. First, we assume the gait sequence; single stance phase followed by heel-strike phase followed by single stance phase and so on. Next, we draw the free body diagram for Ranger model shown in figure 20 b). Finally, we use angular momentum balance about the foot contact points to derive the equations of motion.

Alternately, the equations of motion can be obtained from Ranger's equations presented earlier. We ignore the equation for angular momentum balance of the swing leg about the hip. And in all other equations we set the hip angle to be constant (no acceleration, no discontinuity in velocity at the collision). The equation of single stance are obtained from equation 11 with $i = P_1$. The

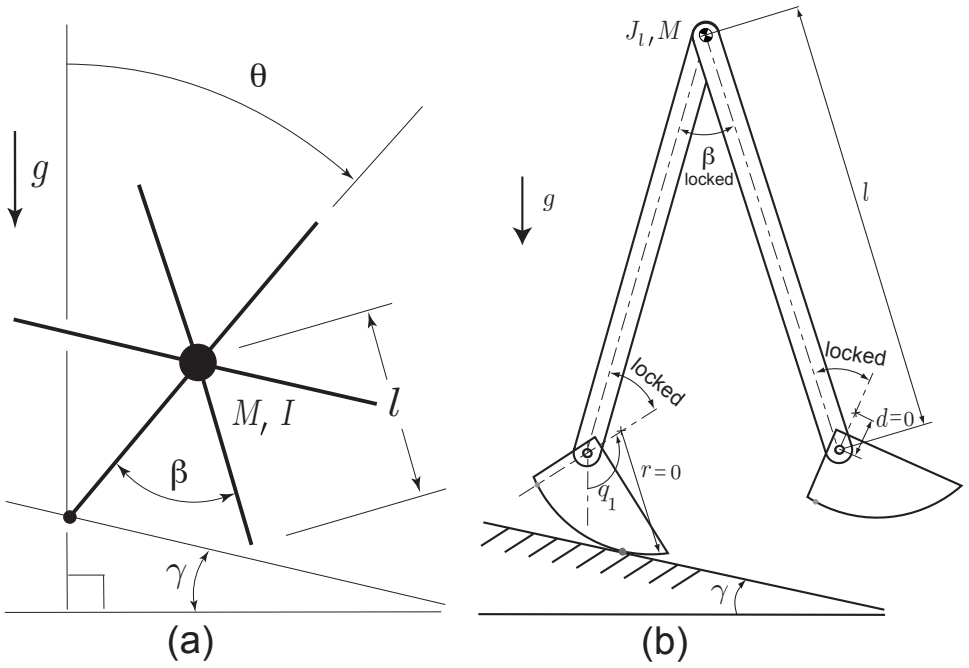


Figure 20: **2-D rimless wheel. (a)** 2-D rimless wheel analyzed by Coleman. (figure source: Coleman's PhD thesis [1]), **(b)** Ranger model simplified to a rimless wheel; The hip angle and ankle angles are locked and the centers of mass of the leg are put at the hip.

equation for heel-strike is obtained from equations 21 and 24.

| | | |
|---|---|---|
| Single stance (continuous): | $\dot{\overrightarrow{H}}_{/P_1} = \overrightarrow{M}_{/P_1}$ | |
| Heel-strike (instantaneous): | $q_1 = r_1 - r_3$ | $q_3 = -r_3$ |
| | $\overrightarrow{H}^{+}_{/P_1} = \overrightarrow{H}^{-}_{/P_2'}$ | |

Table 1 gives the parameters of the 2-D rimless wheel Ranger model.

**Comparison of fixed points.** We analyzed the simplified Ranger model with the parameters given in table 1 and using the recipe presented in section 6.1.

The fixed points based on analytical solutions is [2],
$(\theta_1^*, \dot{\theta}_1^*) = (\pi/n, -0.4603411266094583)$. Using an adaptive step integrator (Runge-Kutta 45) with integration tolerance of $10^{-13}$ and a similar root finder accuracy, we calculated the fixed points as $(q_1^*, \dot{q}_1^*) = (0.523598775598278, -0.460341126609482)$ which is accurate to the 13th decimal place.

**Eigenvalues.** The biggest eigenvalue based on analytical solution [2] is $\sigma = 4/9 = 0.44444444444444\bar{4}$. Using fixed point found earlier and using perturbation of $10^{-5}$, we calculated the Jacobian of the linearized map using central difference. We computed the biggest eigenvalues as $\sigma = 0.444444444411274$. Our Ranger-based eigen-value is accurate to 10th decimal place.

(a) Rimless wheel

| Parameter | Value |
|---|---|
| $\ell$ | 1 m |
| $r$ | 0 |
| $d$ | 0 |
| $w$ | 0 |
| $c$ | 0 |
| $k_h$ | Not in equations |
| $k_s$ | Not in equations |
| $J_\ell$ | 0.25 kg m$^2$ |
| $m$ | 0 |
| $M$ | 1 kg |
| $g$ | 1 m/s$^2$ |
| $\gamma$ | 0.2 |
| $C_{1FW}$ | 0 |
| $C_{2FW}$ | Not in equations |

(b) Simplest walker

| Parameter | Value |
|---|---|
| $\ell$ | 1 m |
| $r$ | 0 |
| $d$ | 0 |
| $w$ | 0 |
| $c$ | 1 m |
| $k_h$ | Not in equations |
| $k_s$ | Not in equations |
| $J_\ell$ | 0 |
| $m$ | 1 kg |
| $M$ | $10^6$ kg |
| $g$ | 1 m/s$^2$ |
| $\gamma$ | 0.009 |
| $C_{1FW}$ | 0 |
| $C_{2FW}$ | Not in equations |

Table 1: **Reduction of Ranger to simpler cases. (a)** Values of Ranger parameters for model reduction to a 2-D rimless wheel. **(b)** Ranger parameters for model reduction to simplest walker.

### 6.3 Reduction to the 2-D simplest walker

A more stringent comparison is with 'The Simplest Walker' which has a non-locked swing leg. Figure 21 (a) shows the 2-D simplest walker analyzed by Garcia [5]. The simplest walker has a point-mass $M$ at the hip. The legs of length $\ell$ are nearly massless, but there is a point-mass $m \ll M$ at the end of the swing leg. Garcia considered the limiting case, $m/M \to 0$. He found that the non-dimensional equations of motion are functions of a single parameter; the slope of the ramp $\gamma$. Garcia [5] does not report benchmark results for the simplest walker. So, using the equations derived by him and for a slope of $\gamma = 0.009$, we computed the fixed points and eigenvalues. These values are reported here and are used as benchmarks.

**Ranger model reduced to the simplest walker.** To derive the equations of motion we proceed as follows. First, we assume the gait sequence; single stance phase followed by heel-strike phase followed by single stance phase and so on. Next, we draw the free body diagram for Ranger model shown in figure 21 b). Finally, we use angular momentum balance about appropriate points to derive equations of motion.

Alternately, the equations of motion can be obtained from Ranger's equations of motion presented earlier. The equation of single stance are obtained from 11 with $i = P_1, H$. The equation for heel-strike is obtained from equations 21,24 and 26 and 24.

| | | |
|---|---|---|
| Single stance (continuous): | $\dot{\overrightarrow{H}}_{/P_1} = \overrightarrow{M}_{/P_1}$ | $\dot{\overrightarrow{H}}_{/H} = \overrightarrow{M}_{/H}$ |
| Heel-strike (instantaneous): | $q_1 = r_1 - r_3$ | $q_3 = -r_3$ |
| | $\overrightarrow{H}^+_{/P_1} = \overrightarrow{H}^-_{/P_2'}$ | $\overrightarrow{H}^+_{/H} = \overrightarrow{H}^-_{/H'}$ |

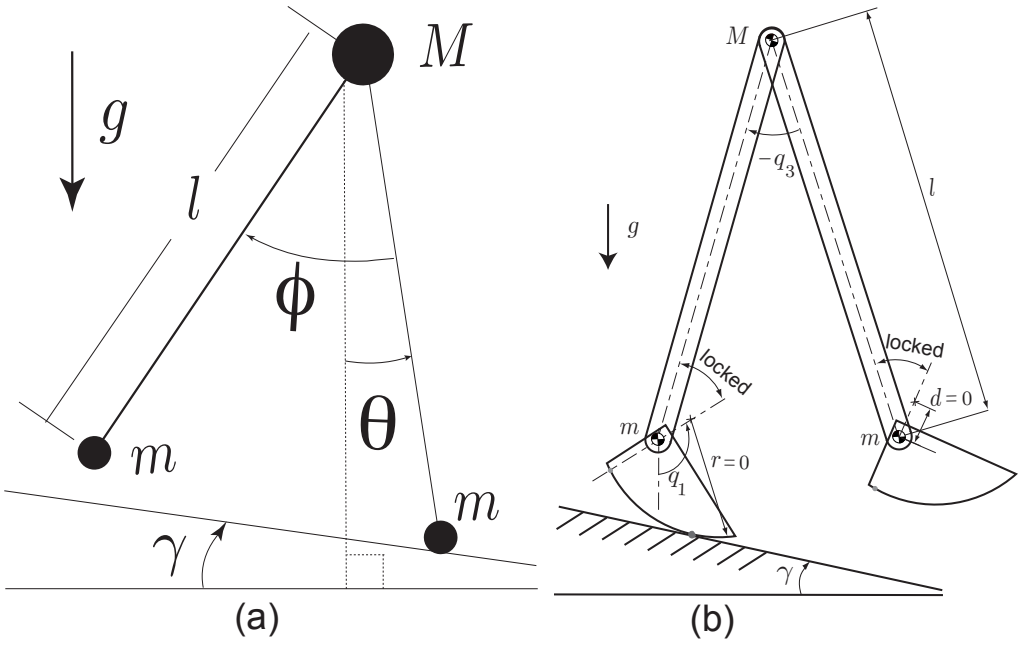


Figure 21: **2-D Simplest walker. (a)** Simplest walker analyzed by Garcia (figure source: Garcia's PhD thesis [6]), **(b)** Ranger model simplified to the simplest walker

Table 1 gives the parameters for the reduced Ranger model to match up with Garcia's simplest walker model.

**Comparison of Ranger's reduction to the simplest walker.** We analyzed the simplified Ranger model using the parameters given in table 1 and using the recipe presented in section 6.1.

**Comparison of fixed points.** First using MATLAB's ODE45 (mixed 4th and 5th order Runge Kutta algorithm) with integration and root finder tolerance of $10^{-13}$, and with Garcia's equations of motion [6], we calculated a fixed point, $(\theta^*,\dot{\theta}^*,\phi^*,\dot{\phi}^*) = (0.200310900544287, -0.199832473004977, 0.400621801088574, -0.015822999948318)$.

Next, using the same tolerances for integration and root finder, but with Ranger's reduced equations of motion, we re-calculated the fixed points. We found the fixed point to be, $(q_1^*,\dot{q}_1^*,q_3^*,\dot{q}_3^*) = (0.200310750572992, -0.199832546623645, 0.400621501145995, -0.015822982402157)$. The fixed points differ in the 6th decimal place. This is consistent with the $m/M = 10^{-6} \neq 0$ that we used.

**Eigenvalues.** Using the fixed point from Garcia's equations of motion and using a perturbation of $10^{-5}$, we calculated the Jacobian of the linearized map using central difference. The non-zero eigenvalues of the linearized map were found to be, $\sigma_1 = -0.190099639087901 + 0.557599274928213i$ and $\sigma_2 = -0.190099639087901 - 0.557599274928213i$.

Using Ranger's reduced equations of motion and the above method to calculate the Jacobian, we got the following non-zero eigenvalues, $\sigma_1 = -0.190106213101483 + 0.557586570259502i$ and $\sigma_2 = -0.190106213101483 - 0.557586570259502i$. The eigenvalues agree to 3 decimals.

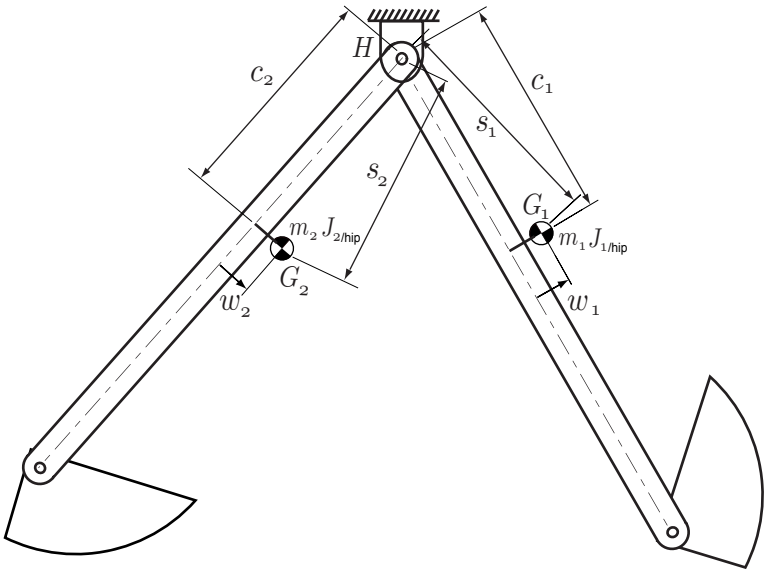


Figure 22: **Dynamic balance of the legs.** The two legs of the robot are dynamically symmetrical ('balanced') if they they have the same distances between the hip and ankle hinges, the same feet shapes and have 3 matching inertial properties. The two masses do not have to be the same, however.

# 7 System identification for mechanical parameters

For numerical simulation we need to estimate the ten parameters shown in figure 1. CAD drawings could have been used to estimate most of these. But because there were many modifications not in the original drawings (rubber feet, glue, tape, etc) we took inertial parameter identification as a separate measurement project. The length parameters (ankle eccentricity $d$, the radius of feet $r$ and leg length $\ell$) are measured with a tape measure. Measurement of the inertial parameters is presented in section 7.2. Measurement of the spring parameters, the hip spring constant $k_h$ and the ankle spring constant $k_s$, is presented below (section 7.3). However, first we show how we dynamically balance the the legs, thereby making the robot symmetrical and simplifying the controller design and simulations.

## 7.1 Dynamic balancing of legs

We would like the legs to be dynamically balanced so that a controller that treats the legs as equal does not lead to a limping gait. Assuming a balanced machine, we can then simplify both the simulations and the controller by only dealing with a single step (rather than a 2-step sequence) as the basic action.

Figure 22 shows that the robot that may have unbalanced legs; i.e., the masses ($m_i$), inertia about hip ($J_{i/hip}$) and the location of the COM ($s_i$) of the two legs may different. Here $i = \ inner$ or *outer*.

**Number of parameters.** In side view the robot has 4 serial links (inner foot, inner leg, outer leg, outer foot). As detached planar rigid objects each of these has 4 inertial parameters: center

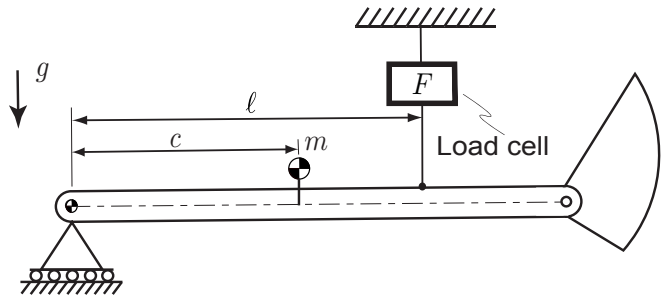


Figure 23: **Experiment to measure the first mass moment $c \cdot m$.** The robot is hinged at the hip joint and held such that the leg axis is perpendicular to gravity. Balance of moments about the hip hinge gives $mc = F\ell$.

of mass position relative to landmarks (e.g., hinges and corners) on the object ($x$ and $y$ positions using an object-based coordinate system); mass $m$; and moment of inertia about the center of mass $J_G$. Thus one could imagine up to 16 inertial parameters in the model. Neglecting the mass of the feet eliminates 8 of these (4 for each foot), reducing the number of inertial parameters to 8.

**Redundancy of parameters.** At every joint we can imagine adding and subtracting point masses. For example we could add a point mass $m_{add}$ to the inner legs at the hip and simultaneously add a negative mass $-m_{add}$ to the outer legs (that such is non-physical does not detract from the argument). No term in any of our equations of motion are affected by this addition and subtraction. Thus no motion nor torque is altered. One could also make the claim by appeal to Lagrange's equations: the expression for the system kinetic and potential energies are unaltered by this addition/subtraction.

We have thus changed the masses of each of the links and the locations of the centers of mass of each of the legs, but we have not changed any of the dynamics of the linkage. [As an aside, with this mass addition/subtraction we have changed the reaction forces transmitted at the hinge, but these do not affect the motions or the joint torques.] Thus, the supposed 8-dimensional parameter space is indifferent to one dimension. That is, there must be a collection of 7 parameters which can predict all coefficients in the governing equations. We claim that the following is such a set of 7 ($c$ and $w$ are local object-referenced $x$ and $y$ coordinates):

$$J_{1/hip},\ J_{2/hip}; \qquad m_1c_1,\ m_2c_2; \qquad m_1w_1,\ m_2w_2; \qquad \text{and} \qquad M_{tot} \equiv m_1 + m_2. \tag{28}$$

Each of these parameters is indifferent to (doesn't change with) the hip-mass addition and subtraction described above. They are also independent in that by appending the list with another single number, for example the mass of one leg, all 8 of the original inertial parameters can be found ($m_i, J_{Gi}, c_i$ and $w_i$). In summary, the 7 mass parameters are: The first mass-moment of the mass distribution of each leg about the hip (2 numbers for each leg), the second polar mass moment of each leg about the hip (polar moment of inertia about the hip), and the total mass of the robot.

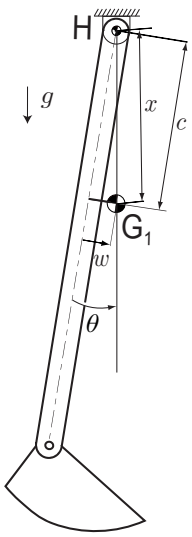


Figure 24: **Experiment to measure the fore-aft offset of the COM.** The robot is hinged at its hip joint $H$ and held in the vertical plane. In equilibrium the center of mass of the leg $G_1$ is directly below the hinge point $H$. The angle $\theta$ can be measured.

## 7.2 Measuring inertial parameters

Not coincidentally, the independent set of 7 parameters are exactly what it is possible to measure without disassembling the robot. Note, for example, that it is not possible to find the mass of one leg ($m_1$ or $m_2$). Nor is it possible to find the distance of the center of mass of one leg from the hip (can't find $s_1$ or $s_2$). We find the 7 parameters as follows:

1) The total mass $M_{tot}$ is found by weighing the robot.

2,3) The first moment of mass along each leg $c \cdot m$ is found for each leg by the experiment shown in figure 23.

4,5) The angle of the radial line from the hip on which the center of mass lies is found by hanging each leg from the hip hinge and measuring the angle $\theta$ of the leg (figure 24. Because $\tan\theta = w/c$ we have for each leg that $wm = \tan\theta\, cm$.

6,7) The inertia of each leg about the hip joint can be found from timing the small oscillations of each leg freely swinging from a hip clamped in place.

$$T_{\text{pend}} = 2\pi\sqrt{\frac{J_{/hip}}{gsm}} \tag{29}$$

where $sm = \sqrt{(cm)^2 + (wm)^2}$. So, for each leg,

$$J_{/hip} = \frac{T_{pend}^2}{4\pi^2} g\sqrt{(cm)^2 + (wm)^2}. \tag{30}$$

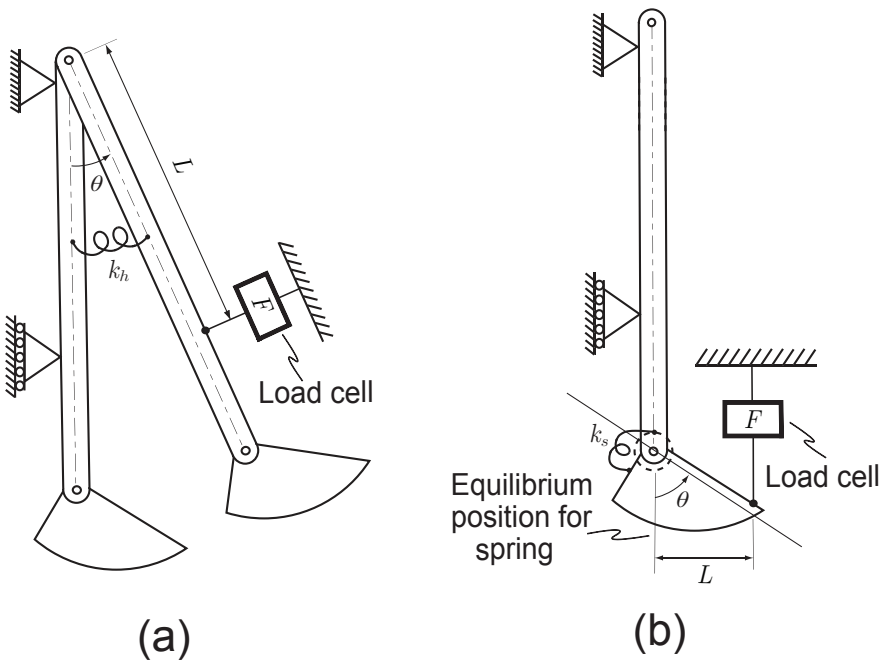


Figure 25: **Measuring the spring constants. a)** The hip spring stiffness (see figure 6 on page 12 and related text) is determined by measuring the force to deflect the leg a given angle. **b)** The ankle spring constants (the elasticities of the cable drive, see figure 3 on page 9 and related text) are measured by locking the ankle motors and measuring the force to deflect the ankles.

Thus all 7 independent inertial properties were measured without robot disassembly.

**Balance of 3 inertial parameters.** With these 7 independent parameters, symmetry of the two legs is achieved by making these 3 matches:

$$\begin{aligned} J_{1/hip} &= J_{2/hip} \\ m_1 c_1 &= m_2 c_2 \\ m_1 w_1 &= m_2 w_2 \end{aligned} \tag{31}$$

As mentioned, although we have four parameters for each leg ($J_{i/G}$, $m_i$, $c_i$ and $w_i$) there are only three conditions that ensure dynamic balance of the two legs.

**Physical balancing.** We made the best balance we could by appropriate placement of the batteries on the outer legs. We chose not to add additional masses. Because we were constrained in our battery attachment points we could not get perfect dynamic balance (i.e., could not achieve equality within measurement accuracy of equation 31).

| Parameter | Value |
|---|---|
| $\ell$ | 0.96 m |
| $r$ | 0.2 m |
| $d$ | 0.11 m |
| $wm$ | 0.0 |
| $cm$ | 0.72 kg m |
| $J_{hip}$ | 0.55 kg m$^2$ |
| $k_h$ | 7.6 N m/rad |
| $k_s$ | 14 N m/rad |

Table 2: Values of robot parameters. These were measured/estimated in bench tests.

### 7.3 Measuring spring constants

**Measuring the hip spring constant.** Figure 25a shows the set up used to estimate the hip spring constant. In this experiment, the robot was placed in the horizontal plane and so gravity does not influence the experiment. The hip springs are designed so they have no torque when the legs are parallel to each other (see figure 6 on page 12). First, we checked this by noting that the spring torque is zero when the legs are parallel. Next, we pulled Ranger's legs with a digital 'fish' scale (a load cell with hooks) and noted the hip angle $\theta$. The hip torsional spring constant is then

$$k_h = FL/\theta.$$

As predicted (see figure 6) the spring is nearly linear in Ranger's operating region (hip angle $\pm 0.5$ rad).

**Measuring the ankle spring constant.** Figure 25b shows the ankle spring-constant measurement. The motor was put in position control (locked at a fixed angle) and the foot deflected with the fish scale (see Figure 25b) and the angle of ankle deflection $\theta$ measured. Thus the ankle spring constant

$$k_s = FL/\theta.$$

Although this measurement includes a parallel contribution from the soft foot return spring (see figure 3), it is dominated by the stiffness of the ankle drive cable (the Achilles tendon).

### 7.4 Summary of parameters estimated

We summarize the various parameters estimated in this section.

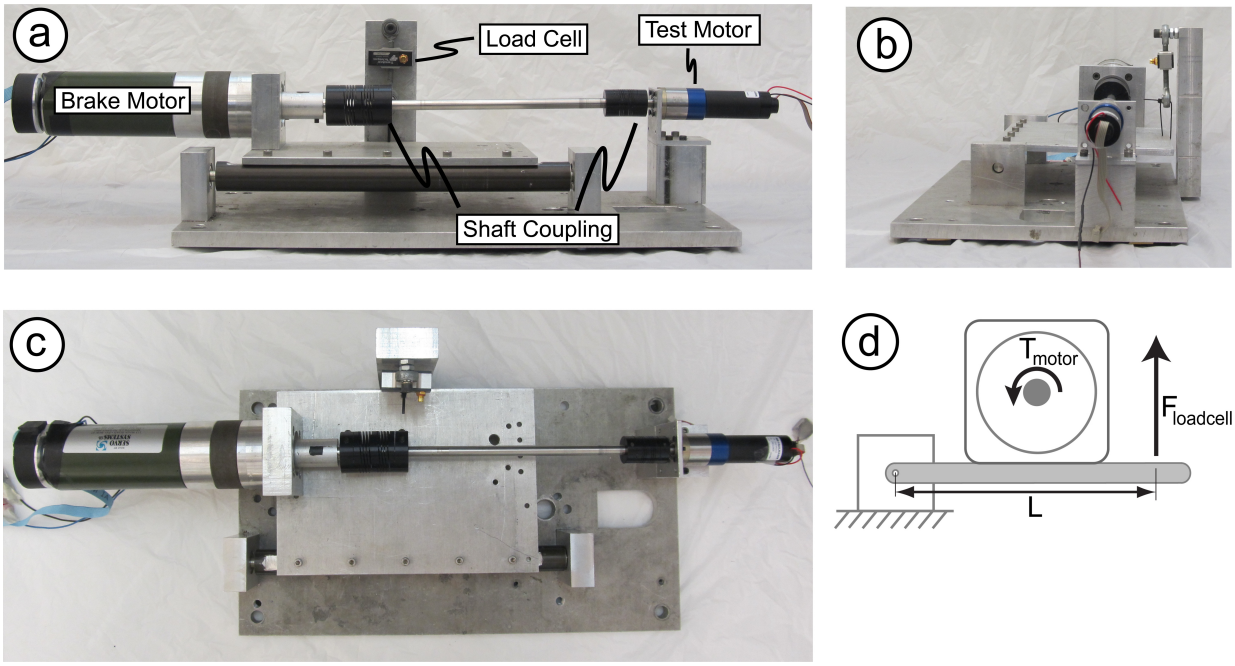


Figure 26: **Cantilever set up used for system identification for DC motors.** A 'brake' motor is mounted on a hinged plate so the torque acting on it can be measured. It is driven by the motor being tested, which is mounted on a solid workbench, via two helical shaft couplings and a slender steel rod. A DC power source is connected to the test motor that maintains the input voltage $V$. By varying the current in the brake motor, varying braking torques can be applied to the test motor. Test motor current $I$, output shaft speed $\omega$, and output shaft torque $T$ are measured and used for system identification. Note that the 'brake' motor can also be powered so that the test motor can be characterized in the negative-work regime (i.e., as a generator).

# 8 Modeling and system identification for motors and gearboxes

Although we use a tight feedback loop on motor current (2 kHz control loop), we choose to not use a tight feedback loop on the motor torque or angular velocity. So we cannot model the motors as pure torque, pure velocity or pure position sources and we need a model for motor torque in terms of motor current and angular velocity. Measuring torque, angular velocity, current and voltage during both positive and negative work on our own bench-test setup, we have found that standard motor models lack two major features: 1) a voltage drop across the brush contacts, and 2) a load dependent, velocity-direction dependent and roughly velocity-magnitude independent frictional torque. Although these phenomena are known, e.g., [4, 7] they are not commonly accounted for, so we review our motor model here.

## 8.1 Motor and gearbox model

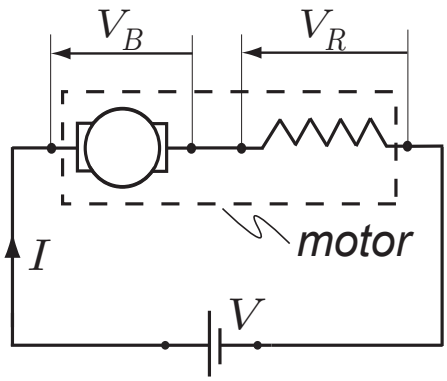


Figure 27: **Schematic of DC motor connected to a DC voltage source.** The DC motor consists of rotating part called the rotor and a resistive part. The sum of voltage drop across the rotor ($V_B$) and resistive part ($V_R$) equals the total voltage supplied by the DC source $V$. The current flowing in the circuit is $I$.

### 8.1.1 Power equation

The electrical power consumption $P$ of the motors is given by the voltage $V$ across the motor times the current $I$ through it. At this point in the modeling we think of a constant Direct Current (DC) source and do not consider losses due to the oscillations in current from the Pulse Width Modulation (PWM) used by our motor controllers. The total electrical power is given as follows.

$$P = VI = V_R I + V_B I \;=\; \overbrace{\left(IR + V_c \operatorname{sgn}(I)\right)}^{V_R} I + \overbrace{GK\omega}^{V_B} I \tag{32}$$

where the signum function $\operatorname{sgn}(I) \equiv I/|I|$. The motor voltage has a contribution from resistance $V_R$ and from "back EMF" $V_B$. The resistance of a real motor has a contribution from the windings, $R$ (the part reported in the motor specification sheets), and a part due to brush contact resistance (which the manufacturer does not mention). We characterize rotation rate using the gearbox output

angular velocity $\omega$ hence $V_B = KG\omega$ is used above (instead of just $K\omega$), where $K$ is the motor constant, $G$ is the down gearing ratio, and $G\omega$ is the motor angular speed.

Here the gear ratio $G$ is known based on manufacturer's specifications and we have to identify the three constants: resistance $R$, contact voltage $V_c$ and torque constant $K$.

### 8.1.2 Torque equation

The output shaft torque ($T$) is given by the ideal-motor output torque after the gearbox $\{G(KI - J_m\dot{\omega_m})\}$ minus the frictional losses in the motor and gear box $T_f$.

$$T = G(\underbrace{KI - J_m \overbrace{G\dot{\omega}}^{\text{motor acceleration}}}_{T_{ideal}}) - T_f(I, \omega) \tag{33}$$

We found that the friction torque $T_f(I, \omega)$ can be reasonably decomposed into a constant friction term and viscous friction term.

$$\begin{aligned} |T_f(I,\omega)| &\leq C_{0s}(I) && \text{if } \omega = 0 \\ T_f(I,\omega) &= C_1\omega + C_{0d}(I)\ \text{sgn}(\omega) && \text{otherwise} \end{aligned} \tag{34}$$

where $C_1$, $C_{0s}$ and $C_{0d}$ are coefficients of viscous, static and dynamic friction respectively and the latter two are assumed to be current dependent.

Next, we found reasonable fit using $C_{0s}(I) = C_{0s} + C'_{0s}|I|$ and $C_{0d}(I) = C_{0d} + C'_{0d}|I|$. We characterize the current dependent part of the constant friction by parameter $\mu$ as follows. We put $C'_{0s} = \mu_s GK|I|$ and $C'_{0d} = \mu_d GK|I|$. Thus our frictional torque becomes,

$$|T_f(I,\omega)| \leq C_{0s} + \mu_s GK|I| \qquad \text{if } \omega = 0 \tag{35}$$

$$T_f(I,\omega) = C_1\omega + C_{0d}\ \text{sgn}(\omega) + \mu_d GK|I|\ \text{sgn}(\omega) \qquad \text{otherwise} \tag{36}$$

In the above equations, the gear ratio $G$ is known based on manufacturer's specifications. The torque constant $K$ is known from the system identification on the power equation presented earlier. Thus in equation 35 and 36, we have to identify the five constants, $C_{0s}$, $C_{0d}$, $\mu_s$, $\mu_d$ and $C_1$.

## 8.2 Cantilever test set-up for data collection

**Details of set-up.** Figure 26 shows the labeled photograph of our experimental set-up. The set-up consists of two motors; a Faulhaber test motor of the type used on the robot and a Maxon brake motor used as the motor load. Each motor can be individually controlled by DC power supplies. The test motor is integrated with a 14:1 gear box. The motors are connected to each other by two helical shaft couplings and a slender steel rod. A Hall-effect current sensor measures the current flowing through the test motor. A rotary encoder measures the angular position of the brake motor. Angular position can be converted to angular speed by finite differencing and low-pass filtering. Using the brake motor gear ratio and its measured angular speed, we can find the speed of the output shaft, the slender steel rod in the set-up. A load cell (Transducer Techniques, MLP-75) measures the test motor torque output. By varying the current flowing through the braking motor, varying braking torques can be applied. Data is recorded by a National Instruments LabVIEW program which is interfaced with the sensors through a data acquisition system.

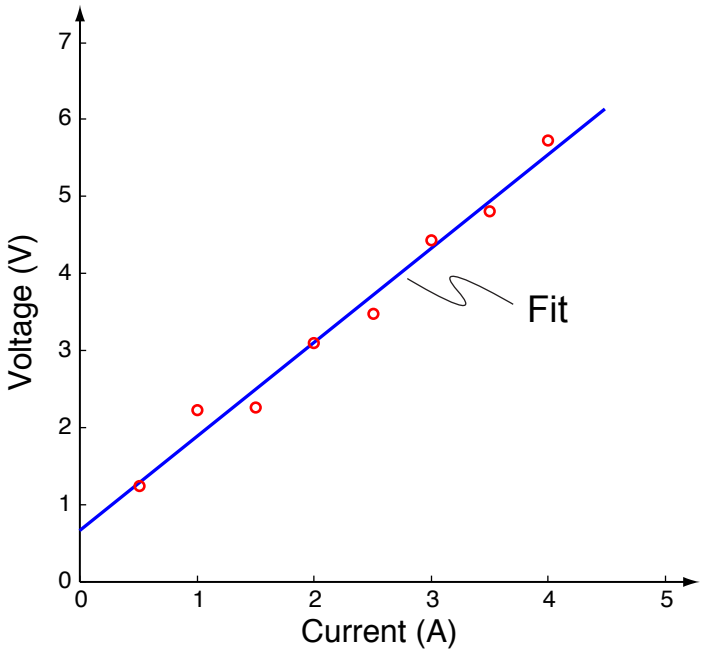


Figure 28: **Curve-fitting for contact voltage and terminal resistance.** Least squares curve-fitting voltage and current data for a stalled motor gives a terminal resistance $R = 1.3\ \Omega$ (slope) and contact voltage drop $V_c = 0.7\ V$ (y-intercept).

**Data collection.** The data collection was done as follows.

- A DC voltage ($V$) was set on the test motor. The braking torque was varied by varying the current to the brake motor. Test motor current ($I$), output shaft speed ($\omega$) and output shaft torque ($T$) reading were noted for increasing and decreasing braking torques.
- A different motor voltage was fixed and the test repeated. The test motor voltages chosen were $-6, -4, -2, 0, 2, 4$ and $6\ V$.

## 8.3 Fit the power equation

In order to fit the power equation **??** we need to find the constants $R$, $V_c$ and $K$. First, we stall the motor and noting the current at various voltages, we fit the constants $V_c$ and $R$. Next, using the data from the cantilever experiment we fit the constant $K$.

### 8.3.1 Fit the resistance ($R$) and contact voltage drop ($V_c$) by stalling the motor

In equation **??**, we first set to identify the resistance $R$ and contact voltage drop $V_c$. We stalled the motor and applied various currents $I$ and measured the output voltage $V$. Putting speed $\omega = 0$ in equation **??** we have $V = IR + V_c\ \mathrm{sgn}(I)$. Using a least squares fit, we found $R = 1.3\ \Omega$ and $V_c = 0.7\ V$. Results of the fit are shown in figure 28.

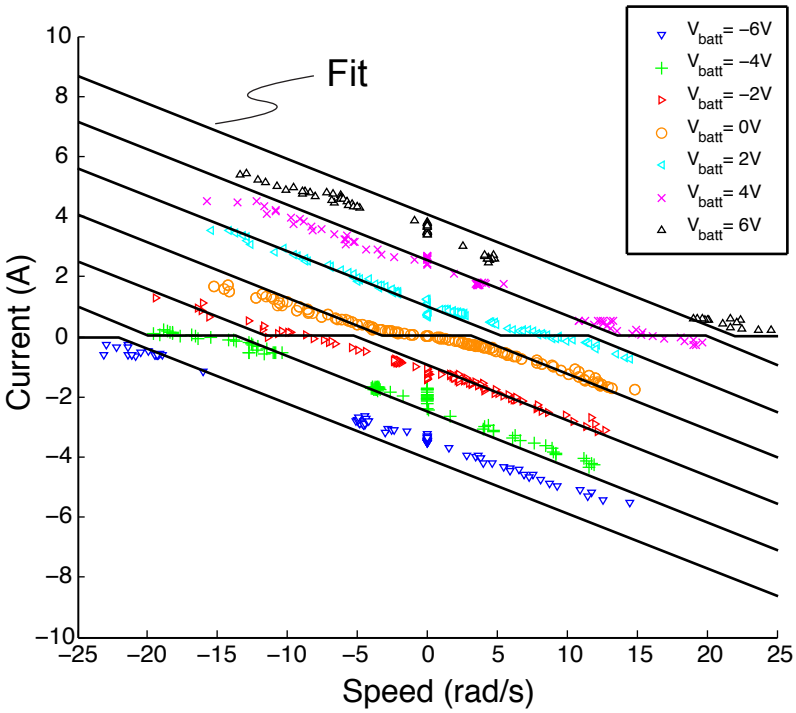


Figure 29: **Curve-fitting for motor torque constant.** Motor torque constant $K = 0.018$ V s/rad was curve-fitted from equation **??** using motor current $I$, output shaft speed $\omega$ and DC voltage $V$ data obtained from the cantilever set-up. In equation **??** we used gear ratio $G = 14$ as per manufacturer's specification. Constants $R = 1.3\ \Omega$ and $V_c = 0.7\ V$ were obtained in an earlier experiment (see figure 28).

### 8.3.2 Fit the the torque constant ($K$) from cantilever experiment

Having fitted the resistance $R$ and contact voltage drop $V_c$ in equation **??**, we only need to fit the torque constant $K$. Using the DC voltage $V$, motor current $I$ and output shaft speed $\omega$ data from the cantilever experiment and knowing that the gear ratio $G = 14$, we fitted the torque constant $K = 0.018$ as shown in figure 29.

## 8.4 Fit the torque equation

In order to fit the torque equations 35 and 36, we have to identify the five constants, $C_{0s}$, $C_{0d}$, $\mu_s$, $\mu_d$ and $C_1$. First, using a series of pulley experiment as shown in figure 30 we identify all the five constants. Next, using the data from the cantilever experiment we check our fit and also point out the necessity of having the $|I|$ function in the frictional torque equations 35 and 36 .

### 8.4.1 Fit the dynamic friction at zero current ($C_{0d}$)

We set the motor in figure 30 in open loop ($I = 0$) and increased the mass $M$ in increments of 5 gram while gently tapping the pulley until it set into slight motion. At about $M = 40$ gram, the mass started to move downward (clockwise rotation of pulley when viewed from the right side). Using the radius of the pulley ($r = 2.5$ cm), mass $M$, gravity $g$ and equation 33 and equation 36,

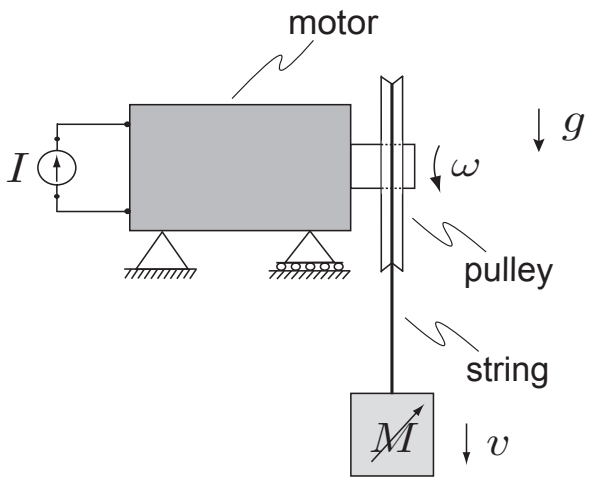


Figure 30: **Measuring friction coefficients.** In some experiments the motor was initially still, in others it is turning at known rate.

we calculated the dynamic friction at zero current to be $T = Mgr = T_f = C_{0d} = 0.01\ N\,m$. An identical value for the dynamic friction was calculated when the mass was hung on the other side of the pulley (counter-clockwise rotation of pulley when viewed from the right side).

### 8.4.2 Fit the static friction at zero current ($C_{0s}$)

We set the motor in figure 30 in open loop ($I = 0$) and increased the mass $M$ in increments of 5 gram till the motor-pulley, with no tapping, until it set into slight motion. At about $M = 45$ gram the mass began to move downwards. Using the radius of the pulley $r$, mass $M$, gravity $g$ and equations 33 and equation 35, we calculated the static friction at zero current to be $T = Mgr = T_f = C_{0s}(0) = 0.01$ N m. The same static friction value was calculated when the test was repeated with the mass hanging on the other side of the pulley.

### 8.4.3 Fit the viscous friction at zero current ($C_1$)

We set the motor in figure 30 in open loop ($I = 0$) and increased the mass to $M_1$ until the motor-pulley set in motion. We noted that speed $\omega_1$. Next we changed the mass to $M_2$ and repeated the experiment and noted the new speed $\omega_2$. From equation 33 and equation 36 we have

$$T_1 = M_1 gr = T_{f1} = C_1\omega_1 + C_{0d}$$
$$T_2 = M_2 gr = T_{f2} = C_1\omega_2 + C_{0d}$$

Subtracted the first equation from the second gives,

$$C_1 = \frac{(M_2 - M_1)gr}{\omega_2 - \omega_1} \tag{37}$$

Using equation 37, we calculated the viscous friction to be $3.3 \times 10^{-3}$ N m s/rad. The same experiment when repeated with the mass hung on the other side of the pulley gave a similar value.

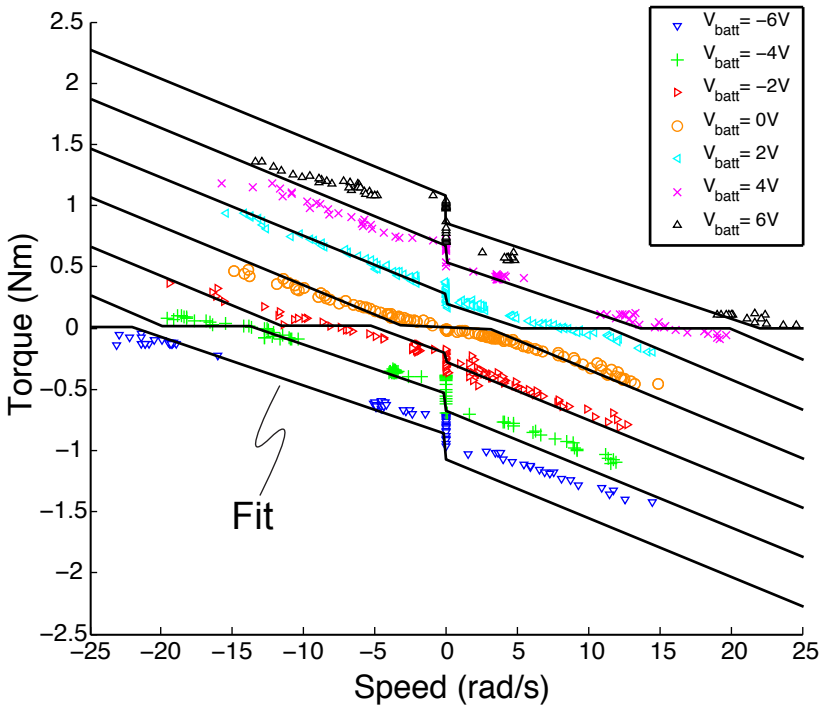


Figure 31: **Checking the friction model.** The frictional torque identified using a series of pulley experiments is checked with data obtained from the cantilever experiment. Various shaped dots are data and the curves are the fit. The positive work regimes are where the torque and angular velocity have the same signs (first and third quadrants). The worst data fits are for high braking torques (lower right and upper left on plots). The discontinuities at speed = 0 are from friction force reversals. The discontinuities near the torque = 0 axis are due to the reversing of the contact voltage drop when the current reverses.

### 8.4.4 Fit the coefficient of dynamic friction ($\mu_d$)

We repeated the experiment in section 8.4.1, except that we set the current to a non-zero value. If $M$ is the load at which the motor-pulley system just starts to move then, from equation 33 and equation 36, just at the onset of motion we have

$$T = Mgr = GKI - \mu_d GK|I| - C_{0d}$$

Solving for $\mu_d$ gives,

$$\mu_d = \frac{1}{\text{sgn}(I)} \left\{ 1 - \frac{Mgr + C_{0d}}{GKI} \right\} \tag{38}$$

We repeated this test for different current values and also by putting the mass on the other side of the pulley. The average value for $\mu_d$ across various tests was found to be 0.1.

| Parameters | Symbol | Expts. (Specs.) |
|---|---|---|
| Terminal resistance ($\Omega$) | $R$ | 1.3 (0.7) |
| Contact voltage drop (V) | $V_c$ | 0.7 |
| Torque constant (N m/A) | $K$ | 0.018 (0.017) |
| Viscous friction (N m s/rad) | $C_1$ | 0 |
| Constant friction (N m) | $C_0 = C_{0s} = C_{0d}$ | 0.01 |
| Current-dependent const. friction | $\mu = \mu_s = \mu_d$ | 0.1 |
| Motor inertia ($kg - m^2$) | $J_m$ | $1.6 \times 10^{-6}$ (from specs) |

Table 3: Comparison between experimental values with manufacturer's specification for the motor model. The static and dynamic constant friction terms have the same value, i.e. $C_{0s} = C_{0d}$ and hence these are replaced with the term $C_0$. Similarly, static and dynamic current dependent friction terms have the same value, i.e. $\mu_s = \mu_d$ and hence these are replaced with the term $\mu$. Note that the measured resistance is almost twice that reported in the specification sheet. Also, the brush-commutator contact voltage drop of the motor is not mentioned in the specification sheet

### 8.4.5 Fit the coefficient of static friction ($\mu_s$)

We repeated the experiment in section 8.4.2, except that we set the current to a non-zero value. Again, if $M$ is the load at which the motor-pulley system just moves when tapped then, from equation 33 and equation 35, just at the onset of motion we have

$$T = Mgr = GKI - \mu_s GK|I| - C_{0s}$$

Solving for $\mu_s$ gives,

$$\mu_s = \frac{1}{\text{sgn}(I)} \left\{ 1 - \frac{Mgr + C_{0s}}{GKI} \right\} \tag{39}$$

We repeated this test for different current values and also by putting the mass on the other side of the pulley. The average value for $\mu_s$ across various tests was found to be 0.1. Finally, we used the value of the constants $C_{0s}$, $C_{0d}$, $C_1$, $\mu_s$ and $\mu_d$ and checked the torque equation with the data obtained from the cantilever experiment. Figure 31 shows our model fit with the torque-speed-voltage data obtained from the cantilever experiment.

## 8.5 Summary of constants for motor model

Table 3 shows all the constants obtained in the motor equation.

All of the bench tests were with a 14:1 gearbox. On the robot we used a 43:1 reduction for the ankles and 66:1 for the hip. These motors and gearboxes should have at least slightly different friction parameter values. However, we did tests of robot leg swing and of foot flip up and down and found that the parameters above gave sufficiently accurate predictions.

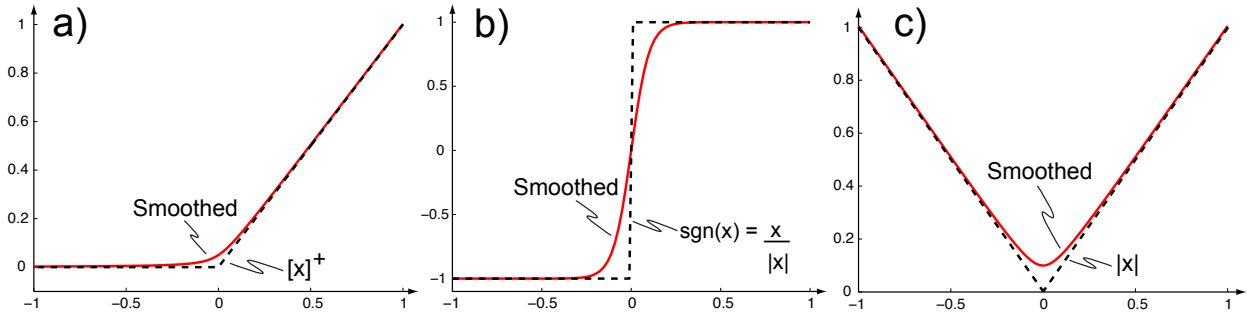


Figure 32: **Smoothings for discontinuous functions.** The smoothings for a unit ramp function, a step function and for the absolute value function (dotted black lines= function, solid red lines = our approximation).

# 9 Smoothings for simulations and optimizations

Because of the reversals of contact voltage at current reversals and of friction force at velocity reversals various terms in the simulations and optimizations are not smooth.

There is some subtlety in the reason for the lack of smoothness in the differential equations to survive to the optimization objective function. If the discontinuities were simply crossed they would not lead to discontinuities in the objective function. But because the optimal trajectories tend to sit on the discontinuities for extended times, not just passing through them, they do survive to the optimizations.

To eliminate the related numerical issues, especially problems with convergence of the optimization software, we smooth such discontinuities.

The smoothings are also needed in our benchmark optimizations based on mechanical power, which, assuming no regeneration, has a discontinuity at $x = 0$.

**Smoothing for $[x]^+$.** Steps (Heaviside functions) in voltage and force only lead to ramps (integral of Heaviside function) in the electrical power (which is integrated to obtain our objective function). So our main concern is with smoothing ramp functions. The unit ramp function $[x]^+$ is zero for negative $x$ and simply $x$ for positive $x$:

$$[x]^+ = \begin{cases} x & \text{if } x > 0 \\ 0 & \text{if } x \leq 0 \end{cases} \tag{40}$$

The function has a kink at $x = 0$. A smooth approximation to $[x]^+$ is given in by [10, 11].

$$[x]^+ \approx \frac{x + \sqrt{x^2 + \epsilon_1^2}}{2} \tag{41}$$

As a rule we use $\epsilon_1 = 0.01$. Figure 32a) compares the smooth approximation with the actual function.

**Smoothing for sgn$(x)$.** When needed we also smooth the signum function $\text{sgn}(x)$:

$$\text{sgn}(x) = \begin{cases} 1 & \text{if } x > 0 \\ 0 & \text{if } x = 0 \\ -1 & \text{if } x < 0 \end{cases} \tag{42}$$

The function is discontinuous at $x = 0$. A smooth, continuous approximation of $\text{sgn}(x)$ is in [3].

$$\text{sgn}(x) \approx \tanh\left(\frac{x}{\epsilon_2}\right) \tag{43}$$

We generally use $\epsilon_2 = 0.01$. Figure 32b) compares the smooth approximation with the actual function.

**Smoothing for $|x|$.** The absolute value function can be defined as:

$$|x| = \begin{cases} x & \text{if } x > 0 \\ -x & \text{if } x \leq 0 \end{cases} \tag{44}$$

This function also has a kink at $x = 0$. A smooth approximation to $|x|$ is given in equation by [10, 11] as

$$|x| \approx \sqrt{x^2 + \epsilon_3^2} \tag{45}$$

Again we generally us $\epsilon_3 = 0.01$. Figure 32c) compares the smooth approximation with the actual function.

# 10 Finite state machine for high level walk control

We use a concurrent, hierarchical finite state machine to code our combined coarse-grid and reflex based discrete controller on the robot.

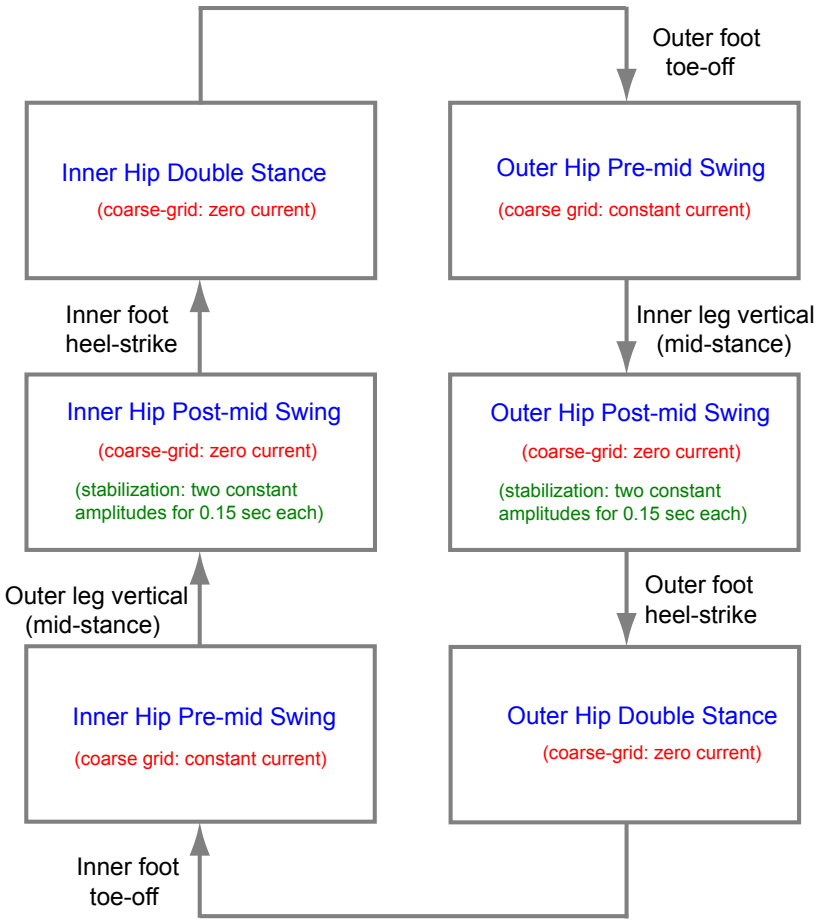


Figure 33: **Hip finite state machine.**

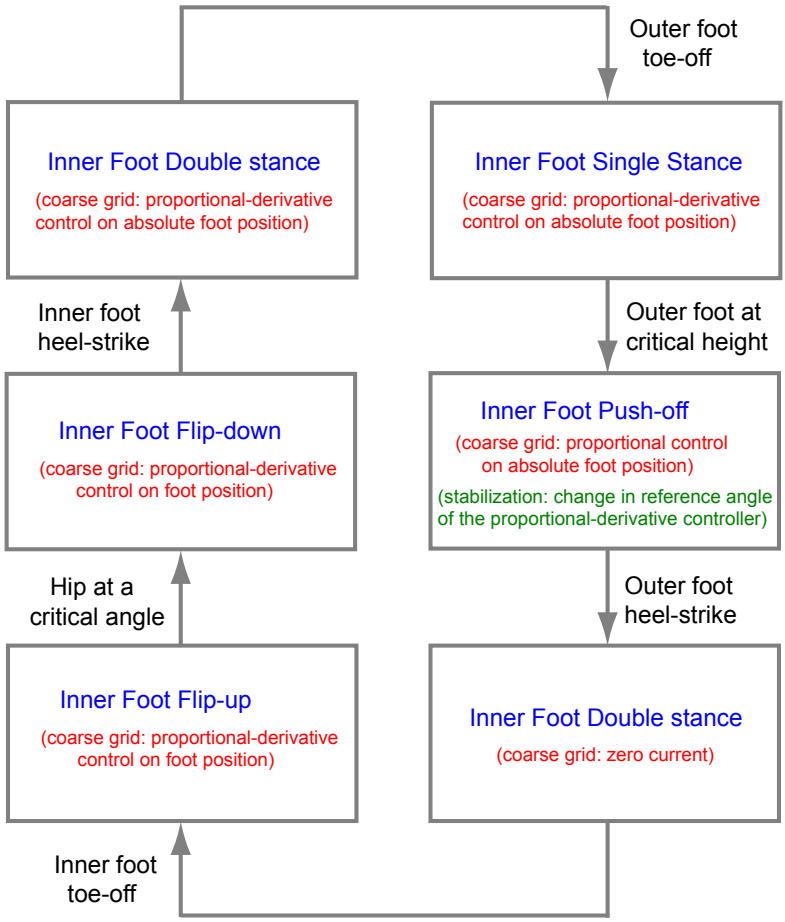


Figure 34: **Inner foot finite state machine.**

## References


[1] M.J. Coleman. *A stability study of a three-dimensional passive-dynamic model of human gait.* PhD thesis, Cornell University, 1998.

[2] M.J. Coleman. Numerical accuracy case studies of two systems with intermittent dynamics a 2d rimless spoked wheel and a 3d passive dynamic model of walking. *Dynamics of Continuous, Discrete and Impulsive Systems Series B: Application and Algorithms*, 16:59–87, 2009.

[3] C. Duan and R. Singh. Dynamics of a 3dof torsional system with a dry friction controlled path. *Journal of sound and vibration*, 289(4-5):657–688, 2006.

[4] P.E. Dupont. The effect of friction on the forward dynamics problem. *The International Journal of Robotics Research*, 12(2):1442–1447, 1993.

[5] M. Garcia, A. Chatterjee, A. Ruina, and M. Coleman. The simplest walking model: Stability, complexity, and scaling. *ASME J. of Biomech. Eng.*, 120:281–288, 1998.

[6] M.S. Garcia. *stability, scaling, and chaos in passive-dynamic gait models.* PhD thesis, Cornell University, 1999.

[7] E. Holm. Contribution to the theory of the contact between a carbon brush and a copper collector ring. *Journal of Applied Physics*, 28:1171–1176, 1957.

[8] T. McGeer. Passive dynamic biped catalogue. In *In Proc. of 2nd International Symposium on Experimental Robotics*, pages 465–490, 1991.

[9] A. Ruina et al. Cornell ranger 2011, 4-legged bipedal robot. [available online]. `http://ruina.tam.cornell.edu/research/topics/locomotion_and_robotics/ranger/Ranger2011/`. Or Google search for: cornell ranger, April 2012.

[10] M. Srinivasan. *Why walk and run: energetic costs and energetic optimality in simple mechanics-based models of a bipedal animal.* PhD thesis, Cornell University, 2006.

[11] K. Taji and M. Miyamoto. A globally convergent smoothing newton method for nonsmooth equations and its application to complementarity problems. *Computational Optimization and Applications*, 22:81–101, 2002.